

Dell DL1000 Appliance User's Guide



Notes, cautions, and warnings



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.



WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Copyright © 2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

2015 - 12

Rev. A01

Contents

1 Introducing your Dell DL1000.....	7
Dell DL1000 core technologies.....	7
Live Recovery.....	7
Universal Recovery.....	7
True Global Deduplication	8
Encryption.....	8
Dell DL1000 data protection features.....	8
Dell DL1000 Core.....	8
Dell DL1000 Smart Agent.....	9
Snapshot process.....	9
Replication — disaster recovery site or service provider.....	9
Recovery.....	10
Recovery-as-a-Service	10
Virtualization and cloud.....	10
Dell DL1000 deployment architecture.....	10
Other information you may need.....	12
2 Working with your DL1000.....	14
Accessing the DL1000 Core Console.....	14
Updating trusted sites in Internet Explorer.....	14
Configuring browsers to remotely access the core console.....	14
Managing licenses	15
Changing a license key	16
Contacting the license portal server	16
Changing the AppAssure language manually.....	16
Changing the OS language during installation.....	17
Managing core settings	17
Changing the core display name	17
Changing the nightly job time	18
Modifying the transfer queue settings	18
Adjusting the client time-out settings	18
Configuring deduplication cache settings	18
Modifying engine settings	19
Modifying deploy settings	20
Modifying database connection settings	20
Managing Events	21
Configuring notification groups	21
Configuring an email server.....	22

Configuring an Email notification template	23
Configuring repetition reduction	24
Configuring Event Retention	24
Managing repositories	24
Viewing repository details.....	25
Checking a repository	25
Managing security	25
Adding an encryption key	26
Editing an encryption key	26
Changing an encryption key passphrase	26
Importing an encryption key	27
Exporting an encryption key	27
Removing an encryption key	27
Managing cloud accounts	27
Adding a cloud account.....	28
Editing a cloud account.....	29
Configuring cloud account settings.....	29
Removing a Cloud account.....	30
Monitoring your DL1000	30
Upgrading your DL1000.....	30
Repairing your DL1000.....	30
Rapid Appliance Self Recovery.....	30

3 Protecting workstations and servers..... 33

About protecting workstations and servers	33
Deploying an Agent (Push Install)	33
Protecting a machine	34
Pausing and resuming protection	36
Deploying the Agent Software when protecting an agent.....	36
Understanding protection schedules	37
Creating custom schedules.....	38
Modifying protection schedules	38
Configuring protected machine settings	39
Viewing and modifying configuration settings	39
Viewing system information for a machine	40
Viewing license information	41
Modifying transfer settings	41
Archiving data.....	43
Creating an archive	43
Importing an archive	45
Archiving to a cloud.....	47
Viewing system diagnostics	47

Viewing machine logs	47
Uploading machine logs.....	47
Canceling operations on a machine	47
Viewing machine status and other details	47
Managing multiple machines	48
Deploying To Multiple Machines	49
Monitoring the deployment of multiple machines	49
Protecting multiple machines.....	50
Monitoring the protection of multiple machines	51
4 Recovering data.....	52
Managing recovery	52
Managing snapshots and recovery points	52
Viewing recovery points	52
Viewing a specific recovery point.....	53
Mounting a recovery point for a Windows machine	54
Dismounting select recovery points	54
Dismounting all recovery points	55
Mounting a recovery point for a Linux machine	55
Removing recovery points	55
Deleting an orphaned recovery point chain.....	56
Forcing a snapshot	56
Restoring data	57
About exporting protected data from Windows machines to virtual machines.....	57
Managing exports.....	58
Exporting backup information from your Windows machine to a virtual machine	59
Exporting Windows data using ESXi export	59
Exporting Windows data using VMware workstation export	61
Exporting Windows data using Hyper-V export	64
Exporting Windows data using Oracle VirtualBox export	67
Restoring Volumes from a Recovery Point	69
Restoring volumes for a Linux machine using the Command Line	72
Launching Bare Metal Restore for Windows machines	73
Roadmap for performing a Bare Metal Restore for a Windows machine	73
Launching a bare metal restore for a Linux machine	78
Installing the screen utility.....	79
Creating bootable partitions on a Linux machine.....	79
5 Replicating recovery points.....	81
Replication.....	81
Roadmap for performing replication	82
Replicating to a self-managed core.....	82

Replicating to a core managed by a third party.....	86
Replicating a new agent	86
Replicating agent data on a machine	87
Setting replication priority for an agent	87
Monitoring replication	88
Managing replication settings	89
Removing replication	89
Removing a protected machine from replication on the source Core.....	90
Removing a protected machine on the target Core.....	90
Removing a target Core from replication.....	90
Removing a source Core from replication.....	90
Recovering replicated data	91
Understanding failover and failback	91
Performing Failover	91
Performing Failback	92
6 Reporting.....	94
About reports	94
About the reports toolbar	94
About compliance reports	94
About errors reports	95
About the Core Summary Report	95
Repositories summary	95
Agents summary	96
Generating a report for a Core or agent	96
About the Central Management Console Core reports	97
Generating a report from the Central Management Console	97
7 Getting help.....	98
Finding documentation and software updates.....	98
Documentation.....	98
Software updates.....	98
Contacting Dell.....	98
Documentation feedback.....	98

Introducing your Dell DL1000

Your Dell DL1000 combines backup and replication into a unified data protection product. It provides reliable application data recovery from your backups to protect virtual machines and physical machines. Your appliance is capable of handling up to terabytes of data with built-in global deduplication, compression, encryption, and replication to specific private or public cloud infrastructure. Server applications and data can be recovered in minutes for data retention (DR) and compliance purposes.

Your DL1000 supports multi-hypervisor environments on VMware vSphere and Microsoft Hyper-V private and public clouds.

Dell DL1000 core technologies

Your appliance combines the following technologies:

- [Live Recovery](#)
- [Universal Recovery](#)
- [True Global Deduplication](#)
- [Encryption](#)

Live Recovery

Live Recovery is instant recovery technology for VMs or servers. It gives you near-continuous access to data volumes on virtual or physical servers.

DL1000 backup and replication technology records concurrent snapshots of multiple VMs or servers, providing near instantaneous data and system protection. You can resume the use of the server by mounting the recovery point without waiting for a full restore to production storage.

Universal Recovery

Universal Recovery provides unlimited machine restoration flexibility. You can restore your backups from physical systems to VMs, VMs to VMs, VMs to physical systems, or physical systems to physical systems, and carry out bare metal restores to dissimilar hardware.

Universal Recovery technology also accelerates cross-platform moves among virtual machines. For example, moving from VMware to Hyper-V or Hyper-V to VMware. It builds in application-level, item-level, and object-level recovery (individual files, folders, email, calendar items, databases, and applications).

True Global Deduplication

True Global Deduplication eliminates redundant or duplicate data by performing incremental block-level backups of the machines.

The typical disk layout of a server consists of the operating system, application, and data. In most environments, the administrators often use a common version of the server and desktop operating system across multiple systems for effective deployment and management. When backup is performed at the block-level across multiple machines, it provides a more granular view of what is in the backup and what is not, irrespective of the source. This data includes the operating system, the applications, and the application data across the environment.



Figure 1. Diagram of True Global Deduplication

Encryption

Your DL1000 provides encryption to protect backups and data-at-rest from unauthorized access and use, ensuring data privacy. The data can be accessed and decrypted using the encryption key. Encryption is performed inline on snapshot data, at line speeds without impacting performance.

Dell DL1000 data protection features

Dell DL1000 Core

The Core is the central component of the DL1000 deployment architecture. The Core stores and manages machine backups and provides services for backup, recovery, retention, replication, archival, and management. The Core is a self-contained network, addressable computer that runs a 64-bit version of Microsoft Windows Server 2012 R2 Foundation and Standard operating systems. The appliance performs target-based inline compression, encryption, and data deduplication of the data received from the agent. The Core then stores the snapshot backups in the repository, which resides on the appliance. Cores are paired for replication.

The repository resides on internal storage within the Core. The Core is managed by accessing the following URL from a JavaScript enabled web browser: **<https://CORENAME:8006/apprecovery/admin>**.

Dell DL1000 Smart Agent

The Smart Agent is installed on the core-protected machine. The Smart Agent tracks the changed blocks on the disk volume and then snaps an image of the changed blocks at a predefined interval of protection. The incremental block-level snapshots' forever approach prevents repeated copying of the same data from the protected machine to the Core.

After the agent is configured, it uses smart technology to track the changed blocks on the protected disk volumes. When the snapshot is ready, it is rapidly transferred to the Core using intelligent multi-threaded, socket-based connections.

Snapshot process

Your DL1000 protection process begins when a base image is transferred from a protected machine to the Core. In this phase, full copy of the machine is transported across the network under normal operation, followed by incremental snapshots forever. The DL1000 Agent for Windows uses Microsoft Volume Shadow copy Service (VSS) to freeze and quiesce application data to disk to capture a file-system-consistent and an application-consistent backup. When a snapshot is created, the VSS writer on the target server prevents content from being written to the disk. During the process of halting of writing content to disk, all disk I/O operations are queued and resume only after the snapshot is complete, while the operations in progress will be completed and all open files will be closed. The process of creating a shadow copy does not significantly affect the performance of the production system.

Your DL1000 uses Microsoft VSS because it has built-in support for all Windows internal technologies such as NTFS, Registry, Active Directory, to flush data to disk before the snapshot. Additionally, other enterprise applications, such as Microsoft Exchange and SQL, use VSS Writer plug-ins to get notified when a snapshot is being prepared and when they have to flush their used database pages to disk to bring the database to a consistent transactional state. The captured data is rapidly transferred and stored on the Core.

Replication — disaster recovery site or service provider

Replication is the process of copying recovery points from an AppAssure core and transmitting them to another AppAssure core in a separate location for disaster recovery. The process requires a paired source-target relationship between two or more cores.

The source core copies the recovery points of selected protected machines, and then asynchronously and continually transmits the incremental snapshot data to the target core at a remote disaster recovery site. You can configure outbound replication to a company-owned data center or remote disaster recovery site (that is, a "self-managed" target core). Or, you can configure outbound replication to a third-party managed service provider (MSP) or cloud provider that hosts off-site backup and disaster recovery services. When replicating to a third-party target core, you can use built-in work flows that let you request connections and receive automatic feedback notifications.

Replication is managed on a per-protected-machine basis. Any machine (or all machines) protected or replicated on a source core can be configured to replicate to a target core.

Replication is self-optimizing with a unique Read-Match-Write (RMW) algorithm that is tightly coupled with deduplication. With RMW replication, the source and target replication service matches keys before transferring data and then replicates only the compressed, encrypted, deduplicated data across the WAN, resulting in a 10x reduction in bandwidth requirements.

Replication begins with seeding: the initial transfer of deduplicated base images and incremental snapshots of the protected machines, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target core using external media. This is typically useful for large sets of data or sites with slow links. The data in the seeding archive is compressed, encrypted and deduplicated. If the total size of the archive is larger than the space available on the removable media, the archive can span across multiple devices based on the available space on the media. During the seeding process, the incremental recovery points replicate to the target site. After the target core consumes the seeding archive, the newly replicated incremental recovery points automatically synchronize.

Recovery

Recovery can be performed in the local site or the replicated remote site. After the deployment is in steady state with local protection and optional replication, the DL1000 Core allows you to perform recovery using Verified Recovery, Universal Recovery, or Live Recovery.

Recovery-as-a-Service

Managed Service Providers (MSPs) can fully leverage DL1000 as a platform for delivering Recovery As A Service (RaaS). RaaS facilitates complete recovery-in-the-cloud by replicating customers' physical and virtual servers. The service provider's cloud are used as virtual machines to support recovery testing or actual recovery operations. Customers wanting to perform recovery-in-the-cloud can configure replication on their protected machines on the local cores to an AppAssure service provider. In the event of a disaster, the MSPs can instantly spin-up virtual machines for the customer.

The DL1000 is not multi-tenant. The MSPs can use the DL1000 at multiple sites and create a multi-tenant environment at their end.

Virtualization and cloud

The DL1000 Core is cloud-ready, which allows you to leverage the compute capacity of the cloud for recovery and archive.

DL1000 can export any protected or replicated machine to licensed versions of VMware or Hyper-V. With continuous exports, the virtual machine is incrementally updated after every snapshot. The incremental updates are fast and provide standby-clones that are ready to be powered up with a click of a button. The supported virtual machine exports are:

- VMware Workstation or Server on a folder
- Direct export to a Vsphere or VMware ESXi host
- Export to Oracle VirtualBox
- Microsoft Hyper-V Server on Windows Server 2008 (x64)
- Microsoft Hyper-V Server on Windows Server 2008 R2
- Microsoft Hyper-V Server on Windows Server 2012 R2

You can now archive your repository data to the cloud using platforms such as Microsoft Azure, Amazon S3, Rackspace Cloud Block Storage, or other OpenStack-based cloud services.

Dell DL1000 deployment architecture

Your DL1000 deployment architecture consists of local and remote components. The remote components may be optional for those environments that do not require leveraging a disaster recovery

site or a managed service provider for off-site recovery. A basic local deployment consists of a backup server called the Core and one or more protected machines known as the agents. The off-site component is enabled using replication that provides full recovery capabilities in the disaster recovery site. The DL1000 Core uses base images and incremental snapshots to compile recovery points of protected agents.

Also, DL1000 is application-aware because it can detect the presence of Microsoft Exchange and SQL and their respective databases and log files. Backups are performed by using application-aware block-level snapshots. DL1000 performs log truncation of the protected Microsoft Exchange server.

The following diagram depicts a simple DL1000 deployment. DL1000 Agents are installed on machines such as a file server, email server, database server, or virtual machines are connected to and protected by a single DL1000 Core, which consists of the central repository. The Dell software License Portal manages license subscriptions, groups and users for the agents and cores in your environment. The License Portal allows users to log in, activate accounts, download software, and deploy agents and cores per your license for your environment.

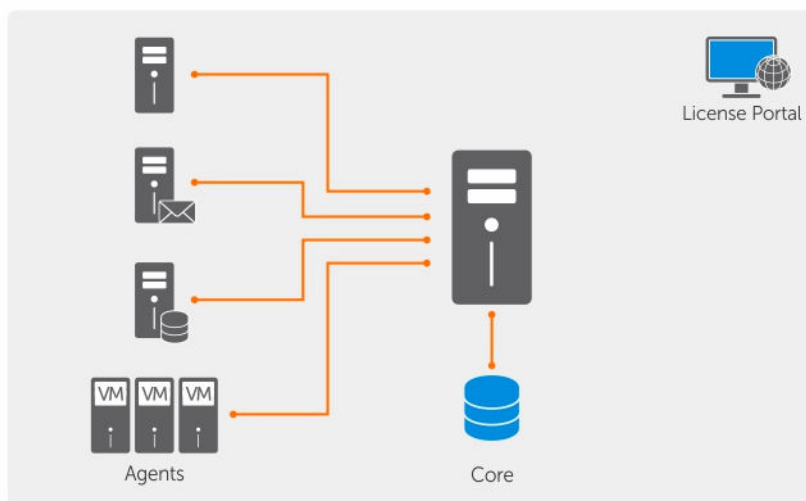


Figure 2. Dell DL1000 Deployment Architecture

You can also deploy multiple DL1000 Cores as shown in the following diagram. A central console manages multiple cores.






Figure 3. DL1000 Multi—Core Deployment Architecture

Other information you may need

-  **NOTE:** For all Dell OpenManage documents, go to dell.com/openmanagemanuals.
-  **NOTE:** Always check for updates on dell.com/support/home and read the updates first because they often supersede information in other documents.
-  **NOTE:** For any documentation related to Dell OpenManage Server Administrator, see dell.com/openmanage/manuals.

Your product documentation includes:

Getting Started Guide	Provides an overview of system features, setting up your system, and technical specifications. This document is also shipped with your system.
System Placemat	Provides information on how to set up the hardware and install the software on your AppAssure solution.
Owner's Manual	Provides information about system features and describes how to troubleshoot the system and install or replace system components.
Deployment Guide	Provides information on hardware deployment and the initial deployment of the appliance.
User's Guide	Provides information about configuring and managing the system.
Release Notes	Provides product information and additional information on the Dell DL1000 Appliance.
Interoperability Guide	Provides information on supported software and hardware for the DL1000 appliance as well as usage considerations, recommendations, and rules.
OpenManage Server Administrator User's Guide	Provides information about using Dell OpenManage Server Administrator to manage your system.

Resource Media

Any media that ships with your system that provides documentation and tools for configuring and managing your system. This includes those pertaining to the operating system, system management software, system updates, and system components that you purchased with your system.

Working with your DL1000

Accessing the DL1000 Core Console

To access the DL1000 Core Console:

1. Update trusted sites in your browser.
2. Configure your browsers to remotely access the DL1000 Core Console. See [Configuring Browsers To Remotely Access The Core Console](#).
3. Perform one of the following to access the DL1000 Core Console:
 - Log on locally to your DL1000 core server, and then double-click the **Core Console** icon.
 - Type one of the following URLs in your web browser:
 - **https://<yourCoreServerName>:8006/apprecovery/admin/core**
 - **https://<yourCoreServerIPAddress>:8006/apprecovery/admin/core**

Updating trusted sites in Internet Explorer

To update the trusted sites in Internet Explorer:

1. Open Internet Explorer.
2. If the **File**, **Edit View**, and other menus are not displayed, press <F10>.
3. Click the **Tools** menu, and select **Internet Options**.
4. In the **Internet Options** window, click the **Security** tab.
5. Click **Trusted Sites** and then click **Sites**.
6. In **Add this website to the zone**, enter **https://[Display Name]**, using the new name you provided for the Display Name.
7. Click **Add**.
8. In **Add this website to the zone**, enter **about:blank**.
9. Click **Add**.
10. Click **Close** and then **OK**.

Configuring browsers to remotely access the core console

To access the Core Console from a remote machine, you need to modify your browser settings.



NOTE: To modify the browser settings, log in to the system as an administrator.



NOTE: Google Chrome uses Microsoft Internet Explorer settings, change Chrome browser settings using Internet Explorer.



NOTE: Ensure that the **Internet Explorer Enhanced Security Configuration** is turned on when you access the Core Web Console either locally or remotely. To turn on the **Internet Explorer Enhanced Security Configuration**:

1. Open **Server Manager**.
2. Select **Local Server IE Enhanced Security Configuration** displayed on the right. Ensure that it is **On**.

To modify browser settings in Internet Explorer and Chrome:

1. Open Internet Explorer.
2. From the **Tools** menu, select **Internet Options, Security** tab.
3. Click **Trusted Sites** and then click **Sites**.
4. Deselect the option **Require server verification (https:) for all sites in the zone**, and then add `http://<hostname or IP Address of the Appliance server hosting the AppAssure Core>` to **Trusted Sites**.
5. Click **Close**, select **Trusted Sites**, and then click **Custom Level**.
6. Scroll to **Miscellaneous** → **Display Mixed Content** and select **Enable**.
7. Scroll to the bottom of the screen to **User Authentication** → **Logon**, and then select **Automatic logon with current user name and password**.
8. Click **OK**, and then select the **Advanced** tab.
9. Scroll to **Multimedia** and select **Play animations in webpages**.
10. Scroll to **Security**, check **Enable Integrated Windows Authentication**, and then click **OK**.

To modify Mozilla Firefox browser settings:

1. In the Firefox address bar, type **about:config**, and then click **I'll be careful, I promise** if prompted.
2. Search for the term **ntlm**.
The search should return at least three results.
3. Double-click **network.automatic-ntlm-auth.trusted-uris** and enter the following setting as appropriate for your machine:
 - For local machines, enter the host name.
 - For remote machines, enter the host name or IP address separated by a comma of the appliance system hosting the AppAssure Core; for example, *IPAddress,host name*.
4. Restart Firefox.

Managing licenses

You can manage your DL1000 licenses directly from the Core Console. From the console, you can change the license key and contact the license server. You can also access the Dell AppAssure License Portal from the **Licensing** page in the Core Console.

The **Licensing** page includes the following information:

- License type
- License status
- Number of machines protected
- Status of last response from the licensing server

- Time of last contact with the licensing server
- Next scheduled attempt of contact with the licensing server
- License constraints

Changing a license key

To change a license key:

1. Navigate to the Core Console, select **Configuration** → **Licensing**.
The **Licensing** page is displayed.
2. From the **License Details** page, click **Change**.
The **Change License Key** dialog box is displayed.
3. In the **Change License Key** dialog box, enter the new license key and then click **OK**.

Contacting the license portal server

The Core Console contacts the portal server to update changes made in the license portal. Communication with the portal server occurs automatically at designated intervals; however, you can initiate communication on demand.

To contact the portal server:

1. Navigate to the Core Console and then click **Configuration** → **Licensing**.
The **Licensing** page is displayed.
2. From the **License Server** option, click **Contact Now**.

Changing the AppAssure language manually

AppAssure allows you to change the language that you had selected while running AppAssure Appliance Configuration Wizard to any of the supported languages.

To change the AppAssure language to the desired language:


1. Launch the registry Editor using `regedit` command.
2. Navigate to **HKEY_LOCAL_MACHINE** → **SOFTWARE** → **AppRecovery** → **Core** → **Localization**.
3. Open **Lcid**.
4. Select **decimal**.
5. Enter the required language value in the `Value` data box, the supported language values are:
 - a. English: 1033
 - b. Brazilian Portuguese: 1046
 - c. Spanish: 1034
 - d. French: 1036
 - e. German: 1031
 - f. Simplified Chinese: 2052
 - g. Japanese: 1041
 - h. Korean: 1042
6. Right-click and restart the services in the given order:
 - a. Windows Management Instrumentation
 - b. SRM Web Service


- c. AppAssure Core
7. Clear the browser cache.
8. Close the browser and restart the core console from the desktop icon.

Changing the OS language during installation

On a running Windows installation, you can use the control panel to select language packs and configure additional international settings.

To change OS language:

 **NOTE:** It is recommended that the OS language and the AppAssure language be set to the same language. otherwise, some messages may be displayed in mixed languages.

 **NOTE:** It is recommended to change the OS language before changing the AppAssure language.


1. On the **Start** page, type `language`, and make sure that the search scope is set to Settings.
2. In the **Results** panel, select **Language**.
3. In the **Change your language preferences** pane, select **Add a language**.
4. Browse or search for the language that you want to install.
For example, select Catalan, and then select Add. Catalan is now added as one of your languages.
5. In the Change your language preferences pane, select **Options** next to the language that you added.
6. If a language pack is available for your language, select `Download and install language pack`.
7. When the language pack is installed, the language is displayed as available to use for the Windows display language.
8. To make this language your display language, move it to the top of your language list.
9. Log out and log in again to Windows for the change to take effect.

Managing core settings

The Core settings are used to define various configuration and performance settings. Most settings are configured for optimal use but you can change the following settings as necessary:

- General
- Nightly Jobs
- Transfer Queue
- Client Timeout Settings
- Deduplication Cache Configuration
- Database Connection Settings

Changing the core display name

 **NOTE:** It is recommended that you select a permanent display name during the initial configuration of your Appliance. If you change it later, you must perform several steps manually to ensure that the new host name takes effect and the appliance functions properly.

To change the Core display name:

1. Navigate to the Core Console, click **Configuration** → **Settings**.
2. In the **General** section, click **Change**.
The **Display Name** dialog box appears.
3. In the **Display Name** text box, enter a new display name for the Core.

4. Click **OK**.

Changing the nightly job time

The Nightly Job option schedules jobs such as rollup, attachability, and truncation for agents protected by the Core.

To adjust the nightly job time:

1. Navigate to the Core Console and select **Configuration → Settings**.
2. In the **Nightly Jobs** section, click **Change**.
The **Nightly Jobs** dialog box appears.
3. In the **Nightly Jobs Time** text box, enter a new start time.
4. Click **OK**.

Modifying the transfer queue settings

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

To modify the transfer queue settings:

1. Navigate to the Core Console, click **Configuration → Settings**.
2. In the **Transfer Queue** section, click **Change**.
The **Transfer Queue** dialog box appears.
3. In the **Maximum Concurrent Transfers** text box, enter a value to update the number of concurrent transfers.
Set a number from 1 to 60. The smaller the number, the lesser the load is on network and other system resources. As the capacity that is processed increases, so does the load on the system.
4. In the **Maximum Retries** text box, enter a value to update the maximum number of retries.
5. Click **OK**.

Adjusting the client time-out settings

Client Timeout Settings specifies the number of seconds or minutes the server waits before it times out when connecting to a client.

To adjust the client time-out settings:

1. Navigate to the Core Console, and click **Configuration → Settings**.
2. In the **Client Timeout Settings Configuration** section, click **Change**.
The **Client Timeout Settings** dialog box appears.
3. In the **Connection Timeout** text box, enter the number of minutes and seconds before a connection time-out occurs.
4. In the **Read/Write Timeout** text box, enter the number of minutes and seconds that you want to lapse before a time-out occurs during a read/write event.
5. Click **OK**.

Configuring deduplication cache settings

Global deduplication reduces the amount of disk storage space required for your backed up data. The Deduplication Volume Manager (DVM) combines a set of storage locations into a single repository. The

deduplication cache holds references to unique blocks. By default, the deduplication cache is 1.5 GB. If the amount of redundant information is so large that the deduplication cache is full, your repository can no longer take full advantage of further deduplication across your repository for newly added data. You can then increase the size of the deduplication cache by changing the deduplication cache configuration in the Core Console.

To configure deduplication cache settings:

1. Navigate to the Core Console, click **Configuration** → **Settings**.
2. In the **Deduplication Cache Configuration** section, click **Change**.
The **Deduplication Cache Configuration** dialog box appears.
3. In the **Primary Cache Location** text box, enter the updated primary cache location.
4. In the **Secondary Cache Location** text box, enter the updated secondary cache location.
5. In the **Metadata Cache Location** text box, enter the updated metadata cache location.
6. Click **OK**.



NOTE: You must restart the Core service for the changes to take effect.

Modifying engine settings

To modify engine settings:

1. Navigate to the Core Console, click **Configuration** → **Settings**.
2. In the **Replay Engine Configuration** section, click **Change**.
The **Replay Engine Configuration** dialog box appears.
3. In the **Replay Engine Configuration** dialog box, specify the **IP address**. Select one of the following:
 - To use the preferred IP address from your TCP/IP, click **Automatically Determined**.
 - To manually enter an IP address, click **Use a specific IP Address**.
4. Enter the configuration information described as follows:

Text Box	Description
Preferable Port	Enter a port number or accept the default setting (8007 is the default port). The port is used to specify the communication channel for the engine.
Admin Group	Enter a new name for the administration group. The default name is BUILTIN\Administrators .
Minimum Async I/O Length	Enter a value or choose the default setting. It describes the minimum asynchronous input/output length. The default setting is 65536.
Receive Buffer Size	Enter an inbound buffer size or accept the default setting. The default setting is 8192.
Send Buffer Size	Enter an outbound buffer size or accept the default setting. The default setting is 8192.
Read Timeout	Enter a read timeout value or choose the default setting. The default setting is 00:00:30.
Write Timeout	Enter a write timeout value or choose the default setting. The default setting is 00:00:30.

5. Select **No Delay**.

6. Click **OK**.

Modifying deploy settings

To modify deploy settings:

1. Navigate to the Core Console and click the **Configuration** tab, and then **Settings**.
2. In the **Deploy Settings** pane, click **Change**.
The **Deploy Settings** dialog box displays.
3. In the **Agent Installer Name** text box, enter the name of the agent executable file. The default is **Agentweb.exe**.
4. In the **Core Address** text box, enter the address for the core.
5. In the **Failed Receive Timeout** text box, enter the number of minutes to wait without activity to timeout.
6. In the **Max Parallel Installs** text box, enter a number for the maximum installations that can be installed in parallel.
7. Select either or both of the following optional settings:
 - Automatic reboot after install
 - Protect After Deploy
8. Click **OK**.

Modifying database connection settings

To modify database connection settings:

1. Navigate to the Core Console, click **Configuration** → **Settings**.
2. In the **Database Connection Settings** section, perform one of the following:
 - To restore the default configuration, click **Restore Default**.
 - To modify the database connection settings, click **Change**.

On clicking change, the **Database Connection Settings** dialog box appears.

3. Enter the settings for modifying the database connection described as follows:

Text Box	Description
Host Name	Enter a host name for the database connection.
Port	Enter a port number for the database connection.
User Name (optional)	Enter a user name for accessing and managing the database connection settings. It is used to specify the log in credentials for accessing the database connection.
Password (optional)	Enter a password for accessing and managing the database connection settings.
Retain event and job history for, days	Enter the number of days to retain the event and job history for the database connection.

4. Click **Test Connection** to verify your settings.
5. Click **Save**.

Managing Events

The Core includes predefined sets of events, which can be used to notify administrators of critical issues on the Core or the backup jobs.

From the **Events** tab, you can manage notification groups, e-mail SMTP settings, Server Settings, Enabled Trace Logs, Cloud Configuration, repetition reduction, and event retention.

The Notification Groups option allows you to manage notification groups, from which you can:

- Specify an event for which you want to generate an alert for the following:
 - Clusters
 - Attachability
 - Jobs
 - Licensing
 - Log Truncation
 - Archive
 - Core Service
 - Export
 - Protection
 - Replication
 - Rollback
- Specify the type of alert (error, warning, and informational).
- Specify to whom and where the alerts are sent. Options include:
 - Email Address
 - Windows Events Logs
 - Syslog Server
- Specify a time threshold for repetition.
- Specify the retention period for all events.

Configuring notification groups

To configure notification groups:

1. From the Core Console, select **Configuration** → **Events**.
2. Click **Add Group**.
The **Add Notification Group** dialog box opens and displays two panels:
 - **Enable Alerts**
 - **Notification Options**

Enabling alerts

Enabling Alerts allows you to define the set of system events that you want to log, create reports, and set alerts.

 **NOTE:** To create alerts for all events, select **All Alerts**.

- To create alerts specific to errors, warnings, informational messages, or a combination, select one of the following:
 - red triangle icon (Error)
 - yellow triangle icon (Warning)
 - blue circle (Information)
 - curved arrow (Restores default)
- To create alerts for specific events click the > symbol next to the relevant group and select the check box to enable the alert.

Configuring notification options

1. In the **Notification Options** panel, specify how to handle the notification process.

The notification options are:


Text Box	Description
Notify by e-mail	Designate the recipients of the email notification. You can enter separate multiple e-mail addresses as well as blind and carbon copies as shown below: <ul style="list-style-type: none"> • To: • CC: • BCC:
Notify by Windows Event Log	Select this option if you want notification of alerts to be reported through the Windows Event Log.
Notify by sys logd	Select this option if you want alerts to be reported through sys logd. Enter the details for the sys logd in the following text boxes: <ul style="list-style-type: none"> • Hostname: • Port:1
Notify by Toast alerts	Select this option if you want the alert to appear as a pop-up in the lower-right corner of your screen.

2. Click **OK**.

The following message is displayed: **The Group name cannot be changed after the creation of the Notification Group. are you sure you want to use this name?.**

- To save the group name, click **Yes**.
- To change the group name, click **No**. Return to the **Notification Options** window, update the group name and other notification group settings, and save your work.

Configuring an email server

 **NOTE:** You must configure notification group settings, including enabling the **Notify by email** option, before sending out email alert messages.

To configure an email server and email notification template:

1. From the Core Console, click **Configuration** → **Events**.
2. In the **Email Settings** pane, click **SMTP server**.
The **SMTP Server Settings** dialog box appears.


3. Enter details for the email server as follows:

Text Box	Description
SMTP Server	Enter the name of the email server to be used by the email notification template. The naming convention includes the host name, domain, and suffix; for example, smtp.gmail.com .
From	Enter a return email address. It is used to specify the return email address for the email notification template; for example, noreply@localhost.com .
Username	Enter a user name for the email server.
Password	Enter a password for accessing the email server.
Port	Enter a port number. It is used to identify the port for the email server; for example, the port 587 for Gmail. The default is 25.
Timeout (seconds)	To specify how long to try a connection before timing out, enter an integer value. It is used to establish the time in seconds when trying to connect to the email server before a time-out occurs. The default is 30 seconds.
TLS	Select this option if the mail server uses a secure connection such as Transport Layer Security (TLS) or Secure Sockets Layer (SSL).

4. Click **Send Test Email** perform the following:
- In the Send Test Email dialog box, enter a destination email address for the test message and click **Send**.
 - If the test message fails, exit the error dialog box and the **Send Test Email** dialog box, and revise your email server configuration settings. Repeat step 4.
 - Click **OK** to confirm.
 - Verify the test email message was sent.
 - Return to the SMTP Server Settings dialog box, click **Save** to close the dialog box and save your settings.

Configuring an Email notification template

To receive email notifications about events you must configure an email server and an email notification template.

 **NOTE:** To receive email alert messages, configure notification group settings and enable the **Notify by email** option.

To configure an email server and email notification template:

- From the Core Console, click **Configuration** → **Events**.
- In the **Email Settings** pane, click **Change**.
The **Edit Email Notification Configuration** dialog box appears.
- Select **Enable email notifications**, and then enter details for the email server as follows:

Text Box	Description
Email Subject	Enter a subject for the email template. It is used to define the subject of the email notification template; for example, <hostname> - <level> <name>.
Email	Enter information for the body of the template that describes the event, when it occurred, and the severity.

4. Click **Send Test Email** perform the following:
 - a. In the Send Test Email dialog box, enter a destination email address for the test message and click **Send**.
 - b. If the test message fails, exit the error dialog box and the Send Test Email dialog box, click **OK** to save the current email template settings, and modify your email server settings, see [Configuring An Email Server And Email Notification Template](#). Ensure you reenter the password for that email account. Save the settings and then return to step 4.
 - c. Click **OK** to confirm.
 - d. Verify the test email message was sent.
 - e. Return to the **Edit Email Notification Configuration** dialog box, click **OK** to close the dialog box and save your settings.

Configuring repetition reduction

To configure repetition reduction:

1. From the Core Console, click **Configuration** → **Events**.
2. From the **Repetition Reduction** section, click **Change**.
The **Enable Repetition Reduction** dialog box appears.
3. Select **Enable Repetition Reduction**.
4. In the **Store events for** text box, enter the number of minutes to store the events for repetition reduction.
5. Click **OK**.

Configuring Event Retention

To configure event retention:

1. From the Core Console, click **Configuration** → **Settings**.
2. Under **Database Connection Settings**, click **change**.
The **Database Connection Settings** dialog box appears.
3. In the **Retain event and job history for** text box, enter the number of days that you want to retain information about events.
For example, you could select 30 days(default).
4. Click **Save**.

Managing repositories

A repository stores the snapshots that are captured from your protected workstations and servers. The repository for your DL1000 is pre-configured. The repository resides on the internal storage of your system.

Key repository concepts and considerations include:

- The repository is based on the AppAssure Scalable Object File System.
- All data stored within a repository is globally deduplicated.
- The Scalable Object File System can deliver scalable I/O performance in tandem with global data deduplication, encryption, and retention management.


Viewing repository details

To view repository details:

1. In the Core Console, click **Configuration** → **Repositories**.
2. Click > next to the **Status** column of the repository for which you want to view the details.
3. Details for the repository include the storage locations and statistics. Storage location details include the metadata path, data path, and size. Statistical information includes:
 - **Deduplication** — Reported as the number of block dedupe hits, block dedupe misses, and block compression rate.
 - **Record I/O** — Consisting of the rate (MB/s), read rate (MB/s), and write rate (MB/s).
 - **Storage Engine** — Include the rate (MB/s) read rate (MB/s), and write rate (MB/s).


Checking a repository

The Core Console can perform a diagnostic check of a repository volume when errors occur. Core errors could be the result of an improper shut down, or a hardware failure.

 **NOTE:** This procedure must only be performed for diagnostic purposes.

To check a repository:

1. Click **Configuration** → **Repositories**.
2. Click the Settings icon next to the Compression Ratio column below the **Actions** button.
3. Click **Check**.
The **Check Repository** dialog box appears.
4. In the **Check Repository** dialog box, click **Check**.

 **NOTE:** When you perform a check, all active tasks associated with this repository will be cancelled. Before the check begins, a message asking you to acknowledge proceeding with the check will display. It is recommended to rebuild recovery points cache. The failure of a check will result in you having to restore the repository from an archive.

Managing security

Your DL1000 provides strong encryption. By doing so, backups of protected machines are inaccessible. Only the user with the encryption key can access and decrypt the data. Encryption does not affect performance. Key security concepts and considerations include:

- Encryption is performed using 256 bit AES in Cipher Block Chaining (CBC) mode that is compliant with SHA-3.
- Deduplication operates within an encryption domain to ensure privacy.
- Encryption is performed without impact on performance.
- You can add, remove, import, export, modify, and delete an encryption key configured on the Core.

Adding an encryption key

To add an encryption key:

1. In the Core Console, click **Configuration** → **Security**.
2. From the **Actions** drop-down menu, click **Add Encryption Key**.
The **Create Encryption Key** dialog box displays.
3. In the **Create Encryption Key** dialog box, enter the details for the key described as follows.

Text Box	Description
Name	Enter a name for the encryption key.
Description	Enter a description of the encryption key. It is used to provide more details for the encryption key.
Passphrase	Enter a passphrase. It is used to control access.
Confirm Passphrase	Re-enter the passphrase. It is used to confirm the passphrase entry.

4. Click **OK**.



CAUTION: It is recommended that you protect the passphrase. If you lose the passphrase, you cannot recover your data.

Editing an encryption key

To edit an encryption key:

1. From the Core Console, click **Configuration** → **Security**.
The **Encryption Keys** screen is displayed.
2. Click > next to the name of the encryption key that you want to edit, and then click **Edit**.
The **Edit Encryption Key** dialog box appears.
3. In the **Edit Encryption Key** dialog box, edit the name or modify the description of the encryption key.
4. Click **OK**.

Changing an encryption key passphrase

To change an encryption key passphrase:

1. From the Core Console, click the **Configuration** → **Security**.
2. Click > next to the name of the encryption key that you want to edit, and then click **Change Passphrase**.
The **Change Passphrase** dialog box appears.
3. In the **Change Passphrase** dialog box, enter the new passphrase for the encryption and then re-enter the passphrase to confirm what you entered.
4. Click **OK**.



CAUTION: It is recommended that you protect the passphrase. If you lose the passphrase, you cannot access the data on the system.

Importing an encryption key

To import an encryption key:

1. From the Core Console, click **Configuration** → **Security**.
2. From the **Actions** drop-down menu, and then click **Import**.
The **Import Key** dialog box appears.
3. In the **Import Key** dialog box, click **Browse** to locate the encryption key that you want to import, and then click **Open**.
4. Click **OK**.

Exporting an encryption key

To export an encryption key:

1. From the Core Console click **Configuration** → **Security**.
2. From the Configuration drop-down menu for the encryption key that you want to export, select **Export**.
The **Export Key** dialog box appears.
3. In the **Export Key** dialog box, click **Save File** to save and store the encryption keys in a secure location.
4. Click **OK**.

Removing an encryption key

To remove an encryption key:

1. From the Core Console, click **Configuration** → **Security**.
2. From the Configuration drop-down menu for the encryption key that you want to remove, select **Remove**.
The **Remove Key** dialog box appears.
3. In the **Remove Key** dialog box, click **OK** to remove the encryption key.



NOTE: Removing an encryption key does decrypt the data.

Managing cloud accounts

Your DL Appliance allows you to backup your data by creating a backup archive of recovery points to a cloud. With your DL Appliance, you can create, edit, and manage your cloud account through a cloud storage provider. You can archive your data to the cloud using Microsoft Azure, Amazon S3, Rackspace Cloud Block Storage, or other OpenStack-based cloud services. See the following topics to manage your cloud accounts:

- [Adding A Cloud Account](#)
- [Editing a Cloud Account](#)
- [Configuring Cloud Account Settings](#)
- [Removing a Cloud Account](#)

Adding a cloud account

Before you can export your archived data to a cloud, add the account for your cloud provider in the Core Console.

To add a cloud account:

1. In the Core Console, click the **Tools** tab.
2. In the left menu, click **Clouds**.
3. On the **Clouds** page, click **Add New Account**.
The **Add New Account** dialog box opens.
4. Select a compatible cloud provider from the **Cloud Type** drop-down list.
5. Enter the details described in the following table based on the cloud type selected in Step 4.

Table 1. Adding a cloud account

Cloud Type	Text Box	Description
Microsoft Azure	Storage Account Name	Enter the name of your Windows Azure storage account.
	Access Key	Enter the access key for your account.
	Display Name	Create a display name for this account in AppAssure; for example, Windows Azure 1.
Amazon S3	Access Key	Enter the access key for your Amazon cloud account.
	Secret Key	Enter the secret key for this account.
	Display Name	Create a display name for this account in AppAssure; for example, Amazon 1.
Powered by OpenStack	User Name	Enter the user name for you OpenStack-based cloud account.
	API Key	Enter the API key for your account.
	Display Name	Create a display name for this account in AppAssure; for example, OpenStack 1.
	Tenant ID	Enter your tenant ID for this account.
	Authentication URL	Enter the authentication URL for this account.

Cloud Type	Text Box	Description
Rackspace Cloud Block Storage	User Name	Enter the user name for your Rackspace cloud account.
	API Key	Enter the API key for this account.
	Display Name	Create a display name for this account in AppAssure; for example, Rackspace 1.

6. Click **Add**.

The dialog box closes, and your account is displayed on the **Clouds** page of the Core Console.

Editing a cloud account

Perform the following steps to edit a cloud account:

1. In the Core Console, click the **Tools** tab.
2. In the left menu, click **Clouds**.
3. Next to the cloud account you want to edit, click the drop-down menu, and then click **Edit**.
The **Edit Account** window opens.
4. Edit the details as necessary, and then click **Save**.



NOTE: You cannot edit the cloud type.

Configuring cloud account settings

The cloud configuration settings let you determine the number of times AppAssure should attempt to connect to your cloud account, and the amount of time spent on an attempt before it times out. To configure the connection settings for your cloud account:

1. In the Core Console, click the **Configuration** tab.
 2. In the left menu, click **Settings**.
 3. On the **Settings** page, scroll down to **Cloud Configuration**.
 4. Click the drop-down menu next to the cloud account you want to configure, and then do one of the following:
 - Click **Edit**.
The **Cloud Configuration** dialog box appears.
1. Use the up and down arrows to edit either of the following options:
 - **Request Timeout:** Displayed in minutes and seconds, it determines the amount of time AppAssure should spend on a single attempt to connect to the cloud account when there is a delay. Connection attempts will cease after the entered amount of time.
 - **Retry Count:** Determines the number of attempts AppAssure should conduct before determining that the cloud account cannot be reached.
 - **Write Buffer Size:** Determines the buffer size reserved for writing archived data to the cloud.
 - **Read Buffer Size:** Determines the block size reserved for reading archived data from the cloud.

2. Click **Next**.
- Click **Reset**. Returns the configuration to the following default settings:
 - **Request Timeout:** 01:30 (minutes and seconds)
 - **Retry Count:** 3 (attempts)

Removing a Cloud account

You can remove a Cloud account to, discontinue your cloud service, or stop using it for a particular Core. To remove a cloud account:

1. On the Core Console, click the **Tools** tab.
2. In the left menu, click **Clouds**.
3. Next to the cloud account you want to edit, click the drop-down menu, and then click **Remove**.
4. In the **Delete Account** window, click **Yes** to confirm that you want to remove the account.
5. If the cloud account is currently in use, a second window asks you if you still want to remove it. Click **Yes** to confirm.



NOTE: Removing an account that is currently in use causes all archive jobs scheduled for this account to fail.

Monitoring your DL1000

You can monitor the status of the DL1000 Appliance subsystems by using the **Appliance** tab **Overall Status** page. The **Overall Status** page displays a status light next to each subsystem, along with a status description indicating the health of the subsystem.

The Overall Status page also provides links to tools that drill down into the details of each subsystem, which can be helpful for troubleshooting warnings or errors. The **System Administrator** link, available for the Appliance Hardware and Storage Hardware subsystems, prompts you to log on to the System Administrator application used for managing hardware. For more information about the System Administrator application, see the *OpenManage Server Administrator User's Guide* on dell.com/support/manuals.

Upgrading your DL1000



NOTE: Dell recommends you download the latest available AppAssure version from the Dell License Activation Portal using the installer.



NOTE: For other software upgrades you will receive a notification to upgrade to the latest version.

Repairing your DL1000




NOTE: Before you begin the repair process, ensure that you stop the Core services.


Rapid Appliance Self Recovery

Rapid Appliance Self Recovery (RASR) is a bare metal restore process where the operating system drives are rebuilt to the default factory image.

To perform the RASR:


 **NOTE:** Dell recommends you to create RASR USB key after you have set up the Appliance. To create RASR USB key, see [Creating the RASR USB Key](#) section .

1. Insert the RASR USB key created.
2. Reboot the appliance through the RASR USB key.
3. Click on **Rapid Appliance Self Recovery**.
A welcome screen is displayed.
4. Click **Next**.
The **Prerequisites** check screen is displayed.

 **NOTE:** Ensure all the hardware, and other prerequisites are checked before performing the RASR.

5. Click **Next**.
The **Recovery Mode Selection** screen is displayed with three options:
 - **System Recovery**
 - **Windows Recovery Wizard**
 - **Factory Reset**
6. Select **Factory Reset** option.
This option will recover the operating system disk from the factory image.
7. Click **Next**.
The **Storage Configuration** screen is displayed.
8. In the **OS Recovery** screen, following warning message is displayed: This operation will recover the operating system. All OS disk data will be overwritten. in a dialog box.
9. Click **Yes**.
The operating system disk starts restoring back to factory reset.
10. Click **Finish**.

Creating the RASR USB key


 **NOTE:** After the initial setup of the software, the **AppAssure Appliance Configuration Wizard** starts automatically. The **Appliance** tab status icon is yellow.

To create a RASR USB key:

1. Navigate to the **Appliance** tab.
2. Using the left pane navigation, select **Appliance** → **Backup**.
Create RASR USB Drive window is displayed.

 **NOTE:** Insert a 16 GB or larger USB key before attempting to create the RASR key.

3. After inserting a 16 GB or larger USB key, click on **Create RASR USB Drive now**.
A **Prerequisite Check** message is displayed.
After the prerequisites are checked **Create the RASR USB Drive** window displays the minimum size required to create the USB drive and **List of Possible target paths**.
4. Select the target and click **Create**.
A warning dialog box is displayed.
5. Click **Yes**.
The RASR USB Drive key is created.


6.  **NOTE:** Make sure to use the Safely Remove USB Drive or the Windows Eject Drive function to prepare the USB key for removal. Otherwise, the content in the USB key may be damaged and the USB key will not work as expected.

Remove the key, label, and store for future use.

Protecting workstations and servers

About protecting workstations and servers

To protect your data using DL1000, add the workstations and servers you want to protect in the Core Console; for example, your Exchange server, SQL Server, or your Linux server.

 **NOTE:** In this chapter, the word *machine* also refers to the AppAssure Agent software installed on that machine.

In the Core Console, you can identify the machine on which an AppAssure Agent software is installed and specify which volumes to protect, define schedules for protection, add extra security measures such as encryption, and more. For more information on how to access the Core Console to protect workstations and servers, see [Protecting A Machine](#).

Deploying an Agent (Push Install)

Your DL1000 lets you deploy the AppAssure Agent Installer to individual Windows machines for protection. Complete the following steps to push the installer to an agent. To deploy agents to multiple machines at the same time, see [Deploying To Multiple Machines](#).

 **NOTE:** Agents must be configured with a security policy that makes remote installation possible.

To deploy an agent:

1. From the Core Console's left navigation area, click **Protected Machines**.
2. Click **Actions** → **Deploy Agent**.
The **Deploy Agent** dialog box appears.
3. In the **Deploy Agent** dialog box, enter the logon settings as described in the following table.

Text Box	Description
Machine	Enter the host name or IP address of the machine that you want to deploy.
Username	Enter the user name to connect to this machine (for example, administrator).
Password	Enter the password to connect to this machine.
Automatic reboot after install	Select to specify whether the Core starts upon the completion of the deployment and installation of the AppAssure Agent Installer.


4. Click **Verify** to validate the credentials you entered.
The **Deploy Agent** dialog box displays a message indicating that validation is being performed.
5. Click **Abort** if you want to cancel the verification process.
After the verification process completes, a message indicating that verification is complete appears.
6. Click **Deploy**.

A message indicating that the deployment has started appears. You can view the progress in the **Events** tab.

7. Click **Show details** to view more information about the status of the agent deployment.
8. Click **OK**.

Protecting a machine

This topic describes how to start protecting the data on a machine that you specify.

 **NOTE:** The machine must have the AppAssure Agent software installed in order to be protected. You can choose to install the AppAssure Agent software prior to this procedure, or you can deploy the software to the agent as you define protection in the **Connection** dialog box. To install the AppAssure Agent software during the process of protecting a machine, see [Deploying The Agent Software When Protecting An Agent](#).

When you add protection, you must specify the name or IP address of the machine to protect and the volumes on that machine to protect as well as define the protection schedule for each volume.


To protect multiple machines at the same time, see [Protecting Multiple Machines](#).

To protect a machine:

1. Reboot the machine on which the AppAssure Agent software is installed, if you haven't already done so.
2. From the Core Console on the core machine, click **Protect** → **Protect Machine** on the button bar. The **Protect Machine Wizard** appears.
3. On the **Welcome** page, select the appropriate installation options:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you do not wish to see the **Welcome** page for the **Protect Machine Wizard** in the future, select the **Skip this Welcome page the next time the wizard opens** option.
4. Click **Next**.
5. On the **Connection** page, enter the information about the machine to which you want to connect as described in the following table.

Text Box	Description
Host	The host name or IP address of the machine that you want to protect.
Port	The port number on which the AppAssure Core communicates with the agent on the machine. The default port number is 8006.
Username	The user name used to connect to this machine; for example, administrator.
Password	The password used to connect to this machine.

6. Click **Next**. If the **Protection** page appears next in the **Protect Machine Wizard**, skip to Step 7.

 **NOTE:** If the **Install Agent** page appears next in the **Protect Machine Wizard**, this indicates that the Agent software is not yet on installed on the designated machine. Click **Next** to install the Agent software. The Agent software must be installed on the machine you want to protect, and that be restarted, before it can back up to the Core. To have the installer reboot the agent machine, select the **After installation, restart the machine automatically (recommended)** option before clicking **Next**.

7. The host name or IP address you specified in the **Connect** dialog box appears in this text field. Optionally, enter a new name for the machine to be displayed in the Core Console.

8. Select the appropriate protection schedule:
 - To use the default protection schedule, in the **Schedule Settings** option, select **Default protection (3 hour snapshots of all volumes)**. With a default protection schedule, the Core will take snapshots of the agent machine once every 3 hours. Snapshots of the agent machine can be taken once every hour (minimum). To change the protection settings at any time after you close the wizard, including choosing which volumes to protect, go to the Summary tab for the specific agent machine.
 - To define a different protection schedule, in the **Schedule Settings** option, select **Custom protection**.
9. Select one of the following:
 - If you selected a Typical configuration from the **Protect Machine Wizard** and specified default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machine you specified.
 - The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) will transfer to the repository on the AppAssure Core following the schedule you defined, unless you specified to initially pause protection.
 - If you selected a Typical configuration for the **Protect Machine Wizard** and specified custom protection, then click **Next** to set up a custom protection schedule. For details on defining a custom protection schedule, see Creating Custom Protection Schedules.
 - If you selected Advanced configuration for the **Protect Machine Wizard**, and default protection, then click **Next** and proceed to Step 12 to see repository and encryption options.
 - If you selected Advanced configuration for the **Protect Machine Wizard** and specified custom protection, then click **Next** and proceed to Step 10 to choose which volumes to protect.
10. On the **Protection Volumes** page, select the volumes on the agent machine that you want to protect. If any volumes are listed that you do not want to include in protection, click in the Check column to clear the selection. Then click **Next**.



NOTE: It is recommended to protect the System Reserved volume and the volume with the operating system (typically the C drive).

11. On the **Protection Schedule** page, define a custom protection schedule.
12. On the **Repository** page, select **Use an existing repository**.
13. Click **Next**.
The **Encryption** page appears.
14. Optionally, to enable encryption, on the **Encryption** page, select **Enable Encryption**.
Encryption key fields appear on the **Encryption** page.



NOTE: If you enable encryption, it will be applied to data for all protected volumes for this agent machine. You can change the settings later from the **Configuration** tab in the Core Console.



CAUTION: AppAssure uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. While using encryption is optional, Dell highly recommends that you establish an encryption key, and that you protect the passphrase you define. Store the passphrase in a secure location as it is critical for data recovery. Without a passphrase, data recovery is not possible.

15. Enter the information as described in the following table to add an encryption key for the Core.

Text Box	Description
Name	Enter a name for the encryption key.
Description	Enter a description to provide additional details for the encryption key.

Text Box	Description
Passphrase	Enter the passphrase used to control access.
Confirm Passphrase	Re-enter the passphrase you just entered.

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

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

Pausing and resuming protection

When you pause protection, you temporarily stop all transfers of data from the current machine.

To pause protection:

1. In the Core Console, click the **Protected Machines** drop-down menu in the left navigation area.
2. Select **Pause Protection** for the machine for which you want to pause protection.
The **Pause Protection** dialog box appears.
3. Select one of the following and click **OK**.
 - If you want to pause protection until you explicitly resume it, select **Pause until resumed**.
 - If you want to pause protection for a specified period, select **Pause for** and then, in the Days, Hours, and Minutes controls, type or select the appropriate pause period as appropriate.



NOTE: To resume protection select **Resume Protection** from the **Protected Machines** drop down menu.

Deploying the Agent Software when protecting an agent

You can download and deploy agents during the process of adding an agent for protection.



NOTE: This procedure is not required if you have already installed the Agent software on a machine that you want to protect.

To deploy agents during the process of adding an agent for protection:

1. Click **Protected Machines** on the left navigation pane.
2. Click **Actions** → **Deploy Agent**.
The **Deploy Agent** dialog box appears.
3. Enter login and protection settings as follows:
 - **Host name** — Specifies the host name or IP address of the machine that you want to protect.
 - **User name** — Specifies the user name used to connect to this machine; for example, administrator.
 - **Password** — Specifies the password used to connect to this machine.
 - **Protect machine after install** — Selecting this option enables AppAssure to take a base snapshot of the data after you add the machine for protection. This option is selected by default. If you deselect this option, then you must force a snapshot manually when you are ready to start data protection.
 - **Display name** — Specifies a name for the machine which appears on the Core Console. The display name could be the same value as the host name.

- **Port** — Specifies the port number on which the Core communications with the Agent on the machine. The default value is 8006.
- **Repository** — Select the repository in which to store data from this agent.



NOTE: You can store data from multiple agents in a single repository.

- **Encryption Key** — Specifies whether encryption should be applied to the data for every volume on this machine to be stored in the repository.



NOTE: You define encryption settings for a repository under the **Configuration** tab in the Core Console.

4. Click **Deploy**.

The **Deploy Agent** dialog box closes. There may be a delay before you see the selected agent appear in the list of protected machines.

Understanding protection schedules

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

Protection schedules are initially defined using the **Protect Machine Wizard** or the **Protect Multiple Machines Wizard**. You can then modify the existing schedule at any time from the Summary tab for a specific agent machine.

AppAssure provides a default protection schedule, with two defined protection periods. The first period is for weekdays (Monday through Friday), with a single time period defined (from 12:00 AM to 11:59 PM). The default interval (the time period between snapshots) is 3 hours. The second period is for weekends (Saturday and Sunday). The default interval for the second period is 3 hours.

When protection is first enabled, the schedule is activated. Thus, using the default settings, regardless of the current time of day, the first backup will occur every 3 hours.

The first backup transfer saved to the Core is called a base image snapshot. All data on all specified volumes (including the operating system, applications, and settings), are saved to the Core. Thereafter, incremental snapshots (smaller backups, consisting only of data changed on the agent since the last backup) are saved to the core regularly, based on the interval defined.

You can create a custom schedule to change the frequency of backups. For example, you can change the interval for the weekday period to 60 minutes, resulting in snapshots every hour. Or you can increase the interval on weekends from 60 minutes to 180 minutes, resulting in snapshots every three hours when traffic is low.

Other options in the **Protection Schedule Wizard** page include for a setting a daily protection time. This results in a single backup daily at the period defined (the default setting is 12:00 PM).

The option to initially pause protection prevents a base image from occurring (and in fact, prevents all backups) until you explicitly resume protection. When you are ready to begin protecting your machines based on the established protection schedule, you must explicitly resume protection.

Creating custom schedules

1. On the **Protection Schedule** page of the **Protect Machine** or **Protect Multiple Machines Wizard**, to change the interval schedule for any period, do the following:
 - a. Select **Periods**.
The existing periods display and can be modified. Editable fields include a start time, end time, and interval (in minutes) for each period.
 - b. Click in the interval field and type an appropriate interval in minutes.
For example, highlight the existing interval and replace it with the value **60** to perform snapshots every 60 minutes in this period.
2. To create a peak and off-peak period for weekdays, change the time range of the weekday period so that it does not include a 24-hour period, set an optimal interval for the peak, select **Take snapshots for the remaining time** and set an off-peak interval, by doing the following:
 - a. Select **Periods**.
The existing periods display and can be modified.
 - b. Click in the **From** box to change the start time for this period.
The **Choose Time** dialog box appears.
 - c. Drag the Hours and Minutes slider controls as appropriate for the desired start time, and then click **Done**. To specify the current time, click **Now**.
 - d. Click in the **To** box to change the end time for this period.
The **Choose Time** dialog box appears.
 - e. Drag the Hours and Minutes slider controls as appropriate for the desired start time, and then click **Done**. To specify the current time, click **Now**.
3. To set a single time of day for a single backup to occur daily, select **Daily protection time** and then enter a time in format HH:MM AM.
4. To define the schedule without beginning backups, select **Initially pause protection**.
Once you pause protection from the wizard, it remains paused until you explicitly resume it. Once you resume protection, backups will occur based on the schedule you established.
5. Click **Finish** or **Next**.


Modifying protection schedules

You can modify the protection schedules for specific volumes on a machine.

To modify protection schedules:

1. In the Core Console, select the machine with a defined protection schedule that you want to change.
The Summary tab displays for the machine.
2. Select the volumes for the protected machine that you want to change, and click **Set a schedule**. To select all volumes at once, click in the checkbox in the header row.
Initially, all volumes share the same protection schedule. Typically, it is good practice to protect, at minimum, the System Reserved volume and the volume with the operating system (typically the C drive).
The **Protection Schedule** dialog box appears.
3. On the **Protection Schedule** dialog box, if you previously created a protection schedule template and want to apply it to this agent, select the template from the drop-down list, and then go to Step 9.
4. If you want to save this new protection schedule as a template, enter a name for the template in the text box.

5. If you want to remove an existing time period from the schedule, clear the check boxes next to each time period option. Options include the following:
 - **Mon - Fri.** This range of time denotes a typical five-day work week.
 - **Sat - Sun.** This range of time denotes a typical weekend.
6. If the weekday start and end times are from 12:00 AM to 11:59 PM, then a single period exists. To change the start or end time of a defined period, do the following:
 - a. Select the appropriate time period.
 - b. Click in the **Start Time** box to change the start time for this period.
 - c. Drag the Hours and Minutes slider controls as appropriate for the desired start time, and then click **Done**. To specify the current time, click **Now**.
 - d. Click in the **End Time** box to change the end time for this period.
The **Choose Time** dialog box appears.
 - e. Drag the Hours and Minutes slider controls as appropriate for the desired start time, and then click **Done**. To specify the current time, click **Now**.
 - f. Change the interval according to your requirements. For example, if defining a peak period, change the interval from 60 minutes to 20 minutes to take snapshots three times hourly.
7. If you defined a period other than 12:00 AM to 11:59 PM in Step 6, then if you want backups to occur in the remaining time ranges, you must add additional periods to define protection by doing the following:
 - a. Click **+ Add period**.
Under the appropriate category (weekdays or weekends), a new time period appears. If the first period started later than 12:00 AM, then AppAssure automatically starts this period at 12:00. Following the above example, this second period starts at 12:00 AM. You may need to adjust hours or minutes for the start and end times.
 - b. Drag the Hours and Minutes slider controls as appropriate for the desired start and end times, as appropriate.
 - c. Change the interval according to your requirements. For example, if defining an off-peak period, change the interval from 60 minutes to 120 minutes to take snapshots every two hours.
8. If needed, continue to create additional periods, setting start and end times and intervals as appropriate.

 **NOTE:** If you want to remove a period you have added, click the **X** to the far right of that period. If you remove a period in error, you can click **Cancel**.
9. When your protection schedule meets your requirements, click **Apply**.
The **Protection Schedule** dialog box closes.


Configuring protected machine settings

After you add protection for machines in AppAssure, you can modify basic machine configuration settings (such as name and host name), protection settings (changing the protection schedule for volumes on the machine, adding or removing volumes, or pausing protection), and more.

Viewing and modifying configuration settings

To view and modify configuration settings:

1. From the Core Console, navigate to the machine that you want to modify.
2. Click the **Configuration** → **Settings**.
3. Click **Change** to modify the machine settings as described in the following table.

Text Box	Description
Display Name	<p>Enter a display name for the machine.</p> <p>A name for this machine to be displayed in the Core Console. By default, this is the host name of the machine. You can change the display name to something more user-friendly if needed.</p>
Host Name	Enter a host name for the machine.
Port	<p>Enter a port number for the machine.</p> <p>The Core uses the default port 8006 to communicate with this machine.</p>
Encryption Key	Edit the encryption key if necessary. Specifies whether encryption is applied to the data for every volume on the machine that is stored in the repository.
Repository	<p>Select a repository for the recovery points. Displays the repository on the Core in which to store the data from this machine.</p> <p> NOTE: This setting can only be changed if there are no recovery points or the previous repository is missing.</p>

Viewing system information for a machine

The Core Console displays all the machines that are being protected .

To view system information for a machine:

1. In the left navigation area of the Core Console, under **Protected Machines**, select the machine to view detailed system information.
2. Click the **Tools** tab.

The System Information tab displays include the following:

- Host Name
- OS Version
- OS Architecture
- Memory (Physical)
- Display Name
- Fully Qualified Domain Name
- Virtual Machine Type (if applicable)

Detailed information about the volumes contained on this machine includes:

- Name
- Device ID
- File System
- Capacity (including Raw, Formatted, and Used)

Other machine information displayed, includes:

- Processors
- Network adapters
- IP addresses associated with this machine

Viewing license information

You can view current license status information for the AppAssure Agent software installed on a machine. To view license information:

1. In the navigation pane, select the machine that you want to view.
2. Click **Configuration** → **Licensing**.
The **Status** screen displays the details about the product licensing.

Modifying transfer settings

You can modify the settings to manage the data transfer processes for a protected machine. The transfer settings described in this section are agent-level settings. To affect transfer at the core level, see [Modifying The Transfer Queue Settings](#).



CAUTION: Changing transfer settings could have dramatic effects on your AppAssure environment. Before modifying transfer settings values, refer to the Transfer Performance Tuning Guide in the Dell AppAssure knowledge base.

There are three types of transfers in DL1000:

Snapshots	The transfer that backs up the data on your protected machine.
VM Export	A type of transfer that creates a virtual machine with all of the backup information and parameters as specified by the schedule defined for protecting the machine.
Restore	A process that restores backup information on a protected machine.

Data transfer in DL1000 involves the transmission of a volume of data along a network from AppAssure Agent machines to the Core. In the case of replication, transfer also occurs from the originating or source Core to the target Core.

Data transfer can be optimized for your system through certain performance option settings. These settings control data bandwidth usage during the process of backing up agent machines, performing VM export, or performing a rollback. Some factors that affect data transfer performance are:

- Number of concurrent agent data transfers
- Number of concurrent data streams
- Amount of data change on disk
- Available network bandwidth
- Repository disk subsystem performance
- Amount of memory available for data buffering






You can adjust the performance options to best support your business needs and fine-tune the performance based on your environment.

To modify transfer settings:

1. In the Core Console, navigate to the machine you want to modify.
2. Click the **Configuration** tab, and then click **Transfer Settings**.
The current **Transfer Settings** page appear.
3. On the **Transfer Settings** page, click **Change**.

The **Transfer Settings** dialog box appears.

4. Enter the **Transfer Settings** options for the machine as described in the following table.

Text Box	Description
Priority	<p>Sets the transfer priority between protected machines. Enables you to assign priority by comparison with other protected machines. Select a number from 1 to 10, with 1 being the highest priority. The default setting establishes a priority of 5.</p> <p> NOTE: Priority is applied to transfers that are in the queue.</p>
Maximum Concurrent Streams	<p>Sets the maximum number of TCP links that are sent to the Core to be processed in parallel per agent.</p> <p> NOTE: Dell recommends setting this value to 8. If you experience dropped packets, try increasing this setting.</p>
Maximum Concurrent Writes	<p>Sets the maximum number of simultaneous disk write actions per agent connection.</p> <p> NOTE: Dell recommends setting this value to the same value that you select for Maximum Concurrent Streams. If you experience packet loss, set this value slightly lower. For example, if Maximum Current Streams is set at 8, set this option to 7.</p>
Maximum Retries	<p>Sets the maximum number of retries for each protected machine, if some of the operations fail to complete.</p>
Maximum Segment Size	<p>Specifies the largest amount of data, in bytes, that a computer can receive in a single TCP segment. The default setting is 4194304.</p> <p> CAUTION: Do not change this option from the default setting.</p>
Maximum Transfer Queue Depth	<p>Specifies the number of commands that can be sent concurrently. You can adjust this option to a higher number if your system as a high number of concurrent input/output operations.</p>
Outstanding Reads per Stream	<p>Specifies how many queued read operations will be stored on the back end. This setting helps to control the queuing of agents.</p> <p> NOTE: Dell recommends setting this value to 24.</p>
Excluded Writers	<p>Select a writer if you want to exclude it. Since the writers that appear in the list are specific to the machine that you are configuring, you may not see all the writers listed. Some writers you may see include:</p> <ul style="list-style-type: none">• ASR Writer• BITS Writer• COM+ REGDB Writer• Performance Counters Writer• Registry Writer• Shadow Copy Optimization Writer

Text Box	Description <ul style="list-style-type: none"> • SQLServerWriter • System Writer • Task Scheduler Writer • VSS Metadata Store Writer • WMI Writer
Transfer Data Server Port	Sets the port for transfers. The default setting is 8009.
Transfer Timeout	Specifies in minutes and seconds the amount of time to allow a packet to be static without transfer.
Snapshot Timeout	Specifies in minutes and seconds the maximum time to wait to take a snapshot.
Network Read Timeout	Specifies in minutes and seconds the maximum time to wait for a read connection. If the network read is not performed in that time, the operation is repeated.
Network Write Timeout	Specifies the maximum time in seconds to wait for a write connection. If the network write is not performed in that time, the operation is repeated.

5. Click **OK**.

Archiving data


Retention policies enforce the periods for which backups are stored on short-term (fast and expensive) media. Sometimes certain business and technical requirements mandate extended retention of these backups, but use of fast storage is cost prohibitive. Therefore, this requirement creates a need for long-term (slow and cheap) storage. Businesses often use long-term storage for archiving both compliance and noncompliance data. The archive feature in AppAssure is used to support the extended retention for compliance and noncompliance data. It is also used to seed replication data to a remote replica core.

Creating an archive

To create an archive:



1. In the Core Console, click **Tools** → **Archive** → **Create**.
The **Add Archive Wizard** dialog box appears.
2. On the **Create** page of the **Add Archive Wizard**, select one of the following options from the **Location Type** drop-down list:
 - Local
 - Network
 - Cloud
3. Enter the details for the archive as described in the following table based on the location type you selected in Step 3.

Table 2. Creating an archive

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, d:\work\archive.
Network	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.
	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list.  NOTE: To select a cloud account, you must first add it to the Core Console. See Adding A Cloud Account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is AppAssure-5-Archive-[DATE CREATED]-[TIME CREATED]

4. Click **Next**.
5. On the **Machines** page of the wizard, select which protected machine or machines contains the recovery points you want to archive.
6. Click **Next**.
7. On the **Options** page, enter the information described in the following table.

Text Box	Description
Maximum Size	Large archives of data can be divided into multiple segments. Select the maximum amount of space you want to reserve for creating the archive by doing one of the following:

Text Box	<p>Description</p> <ul style="list-style-type: none"> • Select Entire Target to reserve all available space in the path provided on the destination provided in Step 4. (for example, if the location is D:\work\archive, all of the available space on the D: drive is reserved). • Select the blank text box, use the up and down arrows to enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve. <p> NOTE: Amazon cloud archives are automatically divided into 50 GB segments. Windows Azure cloud archives are automatically divided into 200 GB segments.</p>
Recycle action	<p>Select one of the following recycle action options:</p> <ul style="list-style-type: none"> • Do not reuse: Does not overwrite or clear any existing archived data from the location. If the location is not empty, the archive write fails. • Replace this Core: Overwrites any pre-existing archived data pertaining to this core but leaves the data for other cores intact. • Erase Completely: Clears all archived data from the directory before writing the new archive. • Incremental: Lets you add recovery points to an existing archive. It compares recovery points to avoid duplicating data that already exists in the archive.
Comment	<p>Enter any additional information that is necessary to capture for the archive. The comment will be displayed if you import the archive later.</p>
Use compatible format	<p>Select this option to archive your data in a format that is compatible with previous versions of cores.</p> <p> NOTE: The new format offers better performance; however it is not compatible with older cores.</p>


8. Click **Next**.
9. On the Date Range page, enter the Start Date and Expiration Date of the recovery points to be archived.
 - To enter a time, click on the time shown (default, 8:00 AM) to reveal the slide bars for selecting hours and minutes.
 - To enter a date, click the text box to reveal the calendar, and then click on the preferred day.
10. Click **Finish**.

Importing an archive

To import an archive:

1. In the Core Console, click **Tools** → **Archive** → **Import**.
2. For **Location Type**, select one of the following options from the drop-down list:
 - Local
 - Network
 - Cloud
3. Enter the details for the archive as described in the following table based on the location type you selected in Step 3.

Table 3. Importing an archive

Option	Text Box	Description
Local	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, d:\work\archiveea.
Network	Output location	Enter the location for the output. It is used to define the location path where you want the archive to reside; for example, \\servername\sharename.
	User Name	Enter a user name. It is used to establish logon credentials for the network share.
	Password	Enter a password for the network path. It is used to establish logon credentials for the network share.
Cloud	Account	Select an account from the drop-down list.  NOTE: To select a cloud account, you must first add it to the Core Console. See Adding A Cloud Account .
	Container	Select a container associated with your account from the drop-down menu.
	Folder Name	Enter a name for the folder in which the archived data is to be saved. The default name is AppAssure-5-Archive-[DATE CREATED]-[TIME CREATED]

4. Click **Check File** to validate the existence of the archive to import. The **Restore** dialog box appears.
5. In the **Restore** dialog box, verify the name of the source core.
6. Select the agents to import from the archive.
7. Select the repository.
8. Click **Restore** to import the archive.

Archiving to a cloud

You can archive your data to a cloud by uploading it to a variety of cloud providers directly from the Core Console. Compatible clouds include Windows Azure, Amazon, Rackspace, and any OpenStack-based provider.

To export an archive to a cloud:

- Add your cloud account to the Core Console. For more information see, [Adding A Cloud Account](#).
- Archive your data and export it to your cloud account.
- Retrieve archived data by importing it from the cloud location.

Viewing system diagnostics

In AppAssure, diagnostic information is available for you to view machine log data for any protected machine. Additionally, you can view and upload diagnostic information for the Core.

Viewing machine logs

If you encounter any errors or issues with the machine, it may be useful to view the logs for troubleshooting purposes.

To view machine logs:

1. In the Core Console, click **Tools** → **Diagnostics** → **View Log**.
The **Download Core Log** page appears.
2. Select **Click here to begin the download**
A message will appear alerting you to open or save the file.
3. Choose your preferred method for handling the log file.

Uploading machine logs

1. Navigate to the Core Console, click **Tools** → **Diagnostics** → **Upload Log**.
The **Upload Log** page is displayed.
2. Select **Click here to begin the upload**.
The Events tab displays for you to view the progress of the upload of log information for the core and all protected machines.

Canceling operations on a machine

You can cancel currently executing operations for a machine. You can cancel a current snapshot or cancel all current operations, which includes exports and replications.

To cancel operations on a machine:

1. In the Core Console, select the machine for which you want to cancel operations.
2. In the **Events** expand the event details for the event or operation you want to cancel.
3. Click **Cancel**.

Viewing machine status and other details

To view machine status and other details:

1. In the Core Console, navigate to the protected machine you want to view.

The information about the machine displays on the **Summary** page. The details that display include the following:

- Host name
- Last Snapshot taken
- Next Snapshot scheduled
- Encryption status
- Version number
- Mountability Check status
- Checksum Check status
- Last Log Truncation performed

Detailed information about the volumes contained on this machine also appear and include:

- Name
- File System type
- Space Usage
- Current Schedule
- Next Snapshot
- Total size
- Used Space
- Free space

If SQL Server is installed on the machine, detailed information about the server also appears and includes:

- Online status
- Name
- Install Path
- Version

If Exchange Server is installed on the machine, detailed information about the server and mail stores also appears and includes:

- Version
- Install Path
- Data Path
- Name Exchange Databases Path
- Log File Path
- Log Prefix
- System Path
- MailStore Type

Managing multiple machines

This topic describes the tasks that administrators perform to deploy AppAssure Agent software simultaneously to multiple Windows machines.

To deploy and protect multiple agents, perform the following tasks:

1. Deploy AppAssure to multiple machines.
See [Deploying To Multiple Machines](#).
2. Monitor the activity of the batch deployment.
See [Monitoring The Deployment Of Multiple Machines](#).
3. Protect multiple machines.
See [Protecting Multiple Machines](#).



NOTE: This step can be skipped if you selected the Protect Machine After Install option during deployment.

4. Monitor the activity of the batch protection.
See [Monitoring the Protection of Multiple Machines](#).

Deploying To Multiple Machines

You can simplify the task of deploying the AppAssure Agent software to multiple Windows machines by using the Bulk Deploy feature of AppAssure. You can bulk deploy to:

- Machines on a VMware vCenter/ESXi virtual host
- Machines on an Active Directory domain
- Machines on any other host

The Bulk Deploy feature automatically detects machines on a host and allows you to select those to which you want to deploy. Alternatively, you can manually enter host and machine information.



NOTE: The machines that you are deploying must have access to the internet to download and install bits as AppAssure uses the web version of the AppAssure Agent Installer to deploy the installation components. If access to the internet is not available, you can push the AppAssure Agent installation program from the Core machine. You can download core and agent updates from the License Portal.

Monitoring the deployment of multiple machines

You can view the progress of the deployment of AppAssure Agent software to the machines.


To monitor the deployment of multiple machines:

1. From the Core Console, click **Events** → **Alerts**.
2. Navigate to the AppAssure Core Home tab and then click the **Events** tab.
Alert events appear in the list, showing the time the event initiated and a message. For each successful deployment of the Agent software, you will see an alert indicating that the protected machine has been added.
3. Optionally, click on any link for a protected machine.
The Summary tab for the selected machine appears, showing pertinent information including:
 - The host name of the protected machine
 - The last snapshot, if applicable
 - The scheduled time of the next snapshot, based on the protection schedule you selected
 - Time Remaining

- The encryption key, if any, used for this protected agent
- The version of the Agent software

Protecting multiple machines

After bulk deploying the AppAssure Agent software to the Windows machines, you must protect the machines to protect the data. If you select **Protect Machine After Install** when you deployed the agent, you can skip this procedure.

 **NOTE:** Agent machines must be configured with a security policy that makes remote installation possible.


To protect multiple machines:

1. From the Core Console, click **Protect** → **Bulk Protect**.
The **Protect Multiple Machines Wizard** window appears.
2. Select the appropriate installation option:
 - If you do not need to define a repository or establish encryption, select **Typical**.
 - If you do not wish to see the Welcome page for the Protect Machine Wizard in the future, select **Skip this Welcome page the next time the wizard opens**.
3. Click **Next**.
The **Connection** page appears.
4. Add the machines that you want to protect by clicking one of the following options.
 - Click **Active Directory** to specify machines on an Active Directory domain. Enter credentials as described in the table below and click **Next**.
 - Click **vCenter/ESXi** to specify virtual machines on a vCenter/ESXi virtual host. Enter credentials as described in the table below and click **Next**.

Text Box	Description
Host	The host name or IP address of the Active Directory domain or of the VMware vCenter Server/ESX(i) virtual host.
Username	Enter the username used to connect to this machine; for example, Administrator.
Password	Enter the secure password used to connect to this machine.

- To add the machines manually, select **Add the machines manually**. Click **Next**.
5. On the **Machines** page, to specify machines manually, type the following connection details for each machine on a separate line, and then click **Next**.
hostname::username::password::port
 6. On the **Machines** page, to specify machines identified from an Active Directory domain or from a VMware vCenter/ESX(i) virtual host, select each appropriate machine you want to protect from the list, and then click **Next**.
The system verifies each machine you added automatically and the **Protection** page appears.
 7. On the **Protection** page, select the appropriate protection schedule:
 - To use the default protection schedule, then in the **Schedule Settings** option, select **Default protection (hourly snapshots of all volumes)**.
 - If you want to define a different protection schedule, then in the Schedule Settings option, select **Custom protection** and then click **Next**.
 8. Proceed with your configuration as follows:
 - If you have selected a Typical configuration for the **Protect Multiple Machines Wizard**, and default protection, then click **Finish** to confirm your choices, close the wizard, and protect the machines you specified.

- If you have selected a Typical configuration for the **Protect Multiple Machines Wizard** and specified custom protection, click **Next** and set up a custom schedule.
 - If you have selected Advanced configuration for the Protect Machine Wizard, click **Next** and proceed to Step 9 to see repository and encryption options.
9. On the **Repository** page select **Use an existing repository**.
 10. Click **Next**.
The **Encryption** page appears.
 11. To enable encryption, on the **Encryption** page, select **Enable Encryption**.
Encryption key fields appear on the **Encryption** page.

 **NOTE:** If you enable encryption, it will be applied to data for all protected volumes for the machines you have specified for protection. You can change the settings later from the **Configuration** tab in the Core Console. For more information about encryption, see [Managing Security](#).
 12. Enter the information as described in the following table to add an encryption key for the Core.

Text Box	Description
Name	Enter a name for the encryption key.
Description	Enter a description to provide additional details for the encryption key.
Passphrase	Enter the passphrase used to control access.
Confirm Passphrase	Re-enter the passphrase you just entered.
 13. Click **Finish** to save and apply your settings.

Monitoring the protection of multiple machines

You can monitor the progress as AppAssure applies the protection policies and schedules to the machines.

To monitor the protection of multiple machines navigate to the Core Console Home tab and click **Events**.

The Events tab displays Tasks, Alerts, and Events. When volumes are transferred, the status, start times, and end times display in the Tasks pane. You can also filter tasks by status (active, waiting, completed and failed).

As each protected machine is added, an alert is logged, indicating if the operation was successful or if errors were logged.

Recovering data

Managing recovery

The AppAssure Core can instantly restore data or recover machines to physical or virtual machines from the recovery points. The recovery points contain agent volume snapshots captured at the block level. These snapshots are application-aware, meaning all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot. Using application-aware snapshots in tandem with Verified Recovery, enables the Core to perform several types of recoveries, including:

- Recovery of files and folders
- Recovery of data volumes, using Live Recovery
- Recovery of data volumes for Microsoft Exchange Server and Microsoft SQL Server, using Live Recovery
- Bare metal restore, using Universal Recovery
- Bare metal restore to dissimilar hardware, using Universal Recovery
- Ad-hoc and continuous export to virtual machines

Managing snapshots and recovery points

A recovery point is a collection of snapshots taken of individual disk volumes and stored in the repository. Snapshots capture and store the state of a disk volume at a given point in time while the applications that generate the data are still in use. In AppAssure, you can force a snapshot, temporarily pause snapshots, and view lists of current recovery points in the repository as well as delete them if needed. Recovery points are used to restore protected machines or to mount to a local file system.

The snapshots that AppAssure captures are captured at the block level and are application aware. This means that all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot.

AppAssure uses a low-level volume filter driver, which attaches to the mounted volumes and then tracks all block-level changes for the next impending snapshot. Microsoft Volume Shadow Services (VSS) is used to facilitate application crash consistent snapshots.

Viewing recovery points

To view recovery points:

1. In the left navigation area of the Core Console, select the machine for which you want to view recovery points, and then click the **Recovery Points** tab.

You can view information about the recovery points for the machine as described in the following table:

Info	Description
------	-------------

Status	Indicates current status of the recovery point.
Encrypted	Indicates if the recovery point is encrypted.
Contents	Lists the volumes included in the recovery point.
Type	Defines a recovery point as either base or differential.
Creation Date	Displays the date when the recovery point was created.
Size	Displays the amount of space that the recovery point consumes in the repository.

Viewing a specific recovery point

To view a specific recovery point:

1. In the left navigation area of the Core Console, select the machine for which you want to view recovery points, and select the **Recovery Points**.
2. Click > next to a recovery point in the list to expand the view.
You can view more detailed information about the contents of the recovery point for the selected machine as well as access a variety of operations that can be performed on the recovery point, described in the following table:

Info	Description
Actions	<p>The Actions menu includes the following operations that you can perform on the selected recovery point:</p> <p>Mount — Select this option to mount the selected recovery point. For more information about mounting a selected recovery point, see Mounting A Recovery Point For A Windows Machine.</p> <p>Export — From the Export option, you can export the selected recovery point to ESXi, VMware workstation, or HyperV.</p> <p>Restore — Select this option to perform a restore from the selected recovery point to a volume you specify.</p>
Contents	<p>The Contents area includes a row for each volume in the expanded recovery point, listing the following information for each volume:</p> <p>The Status indicates current status of the recovery point.</p> <p>Title lists the specific volume in the recovery point.</p> <p>Size displays the amount of space that the recovery point consumes in the repository.</p>

3. Click > next to a volume in the selected recovery point to expand the view.

You can view information about the selected volume in the expanded recovery point as described in the following table:

Text Box	Description
Title	Indicates the specific volume in the recovery point.

Text Box	Description
Raw Capacity	Indicates the amount of raw storage space on the entire volume.
Formatted Capacity	Indicates the amount of storage space on the volume that is available for data after the volume is formatted.
Used Capacity	Indicates the amount of storage space currently used on the volume.

Mounting a recovery point for a Windows machine

In AppAssure, you can mount a recovery point for a Windows machine to access stored data through a local file system.

To mount a recovery point for a Windows machine:

1. From the Core Console, select the machine that you want to mount to a local file system. The **Summary** tab for the selected machine displays.
2. Select the **Recovery Points** tab.
3. In the list of recovery points, click > to expand the recovery point that you want to mount.
4. In the expanded details for that recovery point, click **Mount**. The **Mount Recovery Points** dialog box appears.
5. In the **Mount** dialog box, edit the text boxes for mounting a recovery point as described in the following table:

Text Box	Description
Mount Location: Local Folder	Specify the path used to access the mounted recovery point.
Volume Images	Specify the volume images that you want to mount.
Mount Type	Specify the way to access data for the mounted recovery point: <ul style="list-style-type: none"> • Mount Read-only. • Mount Read-only with previous writes. • Mount Writable.
Create a Windows share for this Mount	Optionally, select the check box to specify whether the mounted recovery point can be shared, and then set access rights to it including the Share name and access groups.

6. Click **Mount** to mount the recovery point.

Dismounting select recovery points

To dismount select recovery points:

1. Navigate to the Core Console , click **Tools** → **Mounts**.
2. On the **Local Mounts** page , next to the mount point for the recovery point you want to dismount click **Dismount**.
3. In the Dismounting the Recovery Point window, click **Yes** to confirm.

Dismounting all recovery points

To dismount all recovery points:

1. Navigate to the Core Console, click **Tools** → **Mounts**.
2. On the **Local Mounts** page, click **Dismount All**.
3. In the **Dismounting the Recovery Point** window, click **Yes** to confirm.

Mounting a recovery point for a Linux machine


Using the **aamount** utility in AppAssure, you can remotely mount a volume from a recovery point as a local volume, on a Linux machine.

1. Create a new directory for mounting the recovery point (for example, you can use the **mkdir** command).
2. Verify the directory exists (for example, by using the **ls** command).
3. Run the AppAssure **aamount** utility as root, or as the super user, for example: **sudo aamount**
4. At the AppAssure mount prompt, enter the following command to list the protected machines. **lm**
5. When prompted, enter the IP address or hostname of your Core server.
6. Enter the logon credentials for the Core server, that is, the user name and password.
A list of the machines that are protected by the AppAssure server will display. Each machine is identified by the following: line item number, host/IP address, and an ID number for the machine. For example: 293cc667-44b4-48ab-91d8-44bc74252a4f
7. Enter the following command to list the recovery points that are available for a specified machine: **lr <line_number_of_machine>**
8. Enter the following command to select and mount the specified recovery point at the specified mount point/path. **m <volume_recovery_point_ID_number> <path>**
9. To verify the mount was successful, enter the following command, which should list the attached remote volume: **l**

Removing recovery points

You can easily remove recovery points for a particular machine from the repository. When you delete recovery points in AppAssure, you can specify one of the following options:

Text Box	Description
Delete All Recovery Points	Removes all recovery points for the selected agent machine from the Repository.
Delete a Range of Recovery Points	Removes all recovery points in a specified range before the current, up to and including the base image, which is all data on the machine as well as all recovery points after the current until the next base image.

 **NOTE:** You cannot recover the recovery points you have deleted.


To remove recovery points:

1. In the left navigation area of the Core Console, select the machine for which you want to view recovery points, and then click the **Recovery Points** tab.
2. Click the **Actions** menu.
3. Select one of the following options:

- To delete all currently stored recovery points, click **Delete All**.
- To delete a set of recovery points in a specific data range, click **Delete Range**. The **Delete** dialog box appears. In the **Delete Range** dialog box, specify the range of recovery points that you want to delete by using a start date and time and an end date and time, and then click **Delete**.


Deleting an orphaned recovery point chain

An orphaned recovery point is an incremental snapshot that is not associated with a base image. Subsequent snapshots continue to build onto this recovery point. Without the base image, the resulting recovery points are incomplete and are unlikely to contain the data needed to complete a recovery. These recovery points are considered to be part of the orphaned recovery point chain. If this situation occurs, the best solution is to delete the chain and create a new base image. For more information about forcing a base image, see [Forcing a Snapshot](#).

 **NOTE:** The ability to delete an orphaned recovery chain is not available for replicated recovery points on a target core.

To delete an orphaned recovery point chain:

1. On the Core Console, select the protected machine for which you want to delete the orphaned recovery point chain.
2. Click the **Recovery Points** tab.
3. Under **Recovery Points**, expand the orphaned recovery point.
This recovery point is labeled in the **Type** column as **Incremental Orphaned**.
4. Next to **Actions**, click **Delete**.
The **Delete Recovery Points** window appears.
5. In the **Delete Recovery Points** window, click **Yes**.

 **CAUTION:** Deleting this recovery point deletes the entire chain of recovery points, including any incremental recovery points that occur before or after it, until the next base image. This operation cannot be undone.

Forcing a snapshot

Forcing a snapshot lets you force a data transfer for the current protected machine. When you force a snapshot, the transfer starts immediately or is added to the queue. Only the data that has changed from a previous recovery point is transferred. If there is no previous recovery point, all data on the protected volumes is transferred, referred to as a base image.

To force a snapshot:

1. In the Core Console, select the machine or cluster with the recovery point for which you want to force a snapshot.
2. Click the **Summary** tab in the **Volumes** section, and then select one of the options described as follows:
 - **Force Snapshot** — Takes an incremental snapshot of data updated since the last snapshot was taken.
 - **Force Base Image** — Takes a complete snapshot of all data on the volumes of the machine.
3. When the notification is displayed in the **Transfer Status** dialog box that the snapshot has been queued, click **OK**.
A progress bar appears next to the machine in the **Machines** tab and displays the progress of the snapshot.

Restoring data

Using AppAssure, you can instantly recover or restore data to your physical machines (for Windows or Linux machines) or to virtual machines from stored recovery points for Windows machines. The topics in this section describe how you can export a specific recovery point for Windows machines to a virtual machine or to roll back a machine to a previous recovery point.

If you have replication set up between two cores (source and target), you can only export data from the target core after the initial replication is complete.

About exporting protected data from Windows machines to virtual machines

AppAssure supports both a one-time export or continuous export (to support virtual standby) of Windows backup information to a virtual machine. Exporting your data to a virtual standby machine provides you with a high availability copy of the data. If a protected machine goes down, you can boot up the virtual machine to then perform recovery.

The following diagram shows a typical deployment for exporting data to a virtual machine.

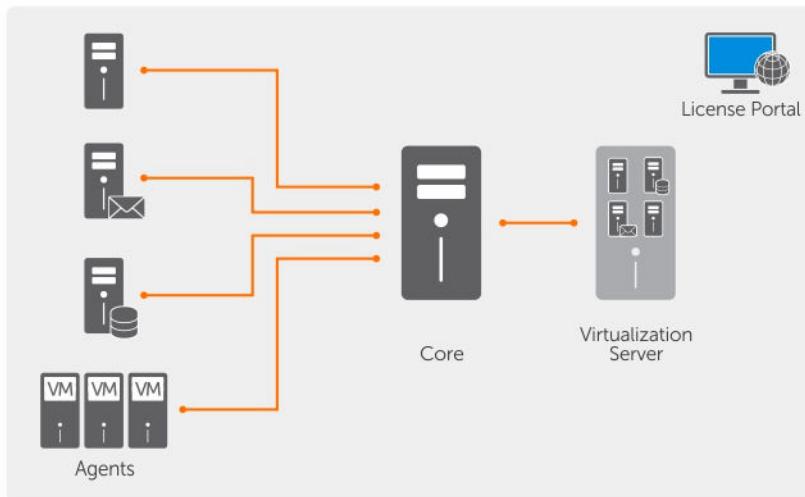




Figure 4. Exporting data to a virtual machine

You create a virtual standby by continuously exporting protected data from your Windows machine to a virtual machine. When you export to a virtual machine, all of the backup data from a recovery point as well as the parameters defined for the protection schedule for your machine will be exported.

You can perform virtual export of recovery points for your protected Windows or Linux machines to VMware, ESXi, Hyper-V, and Oracle VirtualBox.

 **NOTE:** The Appliance tab displays all the virtual machines but only supports the management of Hyper-V and ESXi virtual machines. To manage the other virtual machines use the hypervisor management tools.

 **NOTE:** The virtual machine to which you are exporting must be a licensed version of ESXi, VMWare Workstation, or Hyper-V and not the trial or free versions.


Dynamic and basic volumes support limitations

Dell AppAssure supports taking snapshots of all dynamic and basic volumes. AppAssure also supports exporting simple dynamic volumes that are on a single physical disk. Simple dynamic volumes are not striped, mirrored or spanned volumes.

Dynamic disks (except simple dynamic disks as previously described) are not available for selection in the Export Wizard. Non-simple, dynamic volumes have arbitrary disk geometries that cannot be fully interpreted. AppAssure therefore does not support the export of complex or non-simple dynamic volumes.


Managing exports

On the **Virtual Standby** tab in the Core Console, you can view the status of exports that you have set up, including one-time exports and continuous exports for virtual standby. On this tab, you can manage exports by pausing, stopping, removing exports, or viewing a queue of upcoming exports.

 **NOTE:** Only the Dell DL1000, 3 TB with 2 VMs configuration supports the one-time export and continuous export (virtual standby) capabilities.

1. On the Core Console, navigate to the **Virtual Standby** tab.

On the **Virtual Standby** tab you can view a table of saved export settings, which includes the information described in the following table.

Menu	Description
Status	 NOTE: The status of the virtual standby configuration, is defined by the color of the icon. Green – The Virtual Standby is successfully configured, is active, and not paused. The next Virtual Standby export will be performed after the next snapshot. Yellow – The virtual standby is paused and is still saved by the Core. However, after a new transfer, the export job will not start automatically and there will be no new Virtual Standby exports for this agent.
Machine Name	The name of the source machine.
Destination	The virtual machine and path to which data is being exported.
Export Type	The type of virtual machine platform for the export, such as, ESXi, VMware, Hyper-V, or VirtualBox.
Last Export	The date and time of the last export. If an export has just been added but has not completed, a message will display stating the export has not yet been performed. If an export has failed or was cancelled, a corresponding message also will display.

2. To manage saved export settings, select an export, and then click one of the following:
 - **Pause:** To pause the export.
 - **Resume:** To restart a paused export.


- **Force:** To force a new export. This option could be helpful when virtual standby is paused and then resumed, which means the export job will restart only after a new transfer. If you do not want to wait for the new transfer, you could force an export.
3. To remove an export from the system, click **Remove**. When you remove an export, it is permanently removed from the system and you will not be able to re-start it.
 4. To view details about the active exports currently in queue to be completed, click **Show Export Queue**.

The following table is displayed:

Menu	Description
Machine Name	The name of the source machine.
Destination	The Virtual Standby is successfully configured, is active, and not paused. The next Virtual Standby export will be performed after the next snapshot.
Export Type	The virtual standby is paused and is still saved by the Core. However, after a new transfer, the export job will not start automatically and there will be no new Virtual Standby exports for this agent.
Schedule Type	The type of export as either One-time or Continuous.
Status	The progress of the export, displayed as a percentage in a progress bar.

Exporting backup information from your Windows machine to a virtual machine

You can export data from your Windows machines to a virtual machine (VMware, ESXi, and Hyper-V) by exporting all of the backup information from a recovery point as well as the parameters defined for the protection schedule for your machine.

 **NOTE:** Only the Dell DL1000, 3 TB with 2 VMs configuration supports the one-time export and continuous export (virtual standby) capabilities.

To export Windows backup information to a virtual machine:

1. In the Core Console, click the **Protected Machines** tab.
2. In the list of protected machines, select the machine or cluster with the recovery point for which you want to export.
3. In the **Actions** drop-down menu for that machine, click **Export**, and then select the type of export you want to perform. You can choose from the following options:
 - One-time
 - Virtual Standby

The **Export Wizard** dialog box appears.

Exporting Windows data using ESXi export

In AppAssure, you can choose to export data using ESXi Export by performing a one-time or continuous export.

Performing a one-time ESXi export

To perform a one-time ESXi export:

1. In the Core Console, navigate to the machine you want to export.
2. On the **Summary** tab, click **Actions** → **Export** → **One-time**.
The **Export Wizard** displays on the **Protected Machines** page.
3. Select a machine for export for export and then click **Next**.
4. On the **Recovery Points** page, select the recovery point that you want to export, and then click **Next**.

Defining virtual machine information for performing an ESXi export

To define virtual machine information for performing an ESXi export:

1. On the **Destination** page in the **Export Wizard**, in the **Recover to Virtual machine** drop-down menu, select **ESXi(i)**.
2. Enter the parameters for accessing the virtual machine described as follows:

Text Box	Description
Host name	Enter a name for the host machine.
Port	Enter the port for the host machine. The default port is 443.
User name	Enter the logon credentials for the host machine.
Password	Enter the logon credentials for the host machine.

3. On the **Virtual Machine Options** page, enter the information described in the following table.

Text Box	Description
Resource Pool	Select a resource pool from the drop-down list.
Data Store	Select a data store from the drop-down list.
Virtual Machine Name	Enter a name for the Virtual Machine.
Memory	Specify the memory usage.
Disk Provisioning	Select the type of disk provisioning as either Thin or Thick.
Disk Mapping	Specify the type of disk mapping as either Automatic or Manual.
Version	Select the version of the virtual machine.

4. Click **Next**.
5. On the **Volumes** page, select the volumes you want to export, and then click **Next**.
6. On the **Summary** page, click **Finish** to complete the wizard and start the export.



NOTE: You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Performing a continuous (virtual standby) ESXi export

To perform a continuous (virtual standby) ESXi export:

1. In the Core Console, perform one of the following:

- On the **Virtual Standby** tab, click **Add** to launch the **Export Wizard**. On the **Protected Machines** page of the **Export Wizard**, select the protected machine you want to export, and then click **Next**.
 - Navigate to the machine you want to export, and click **Actions** → **Export** → **Virtual Standby**.
2. On the **Destination** page of the **Export Wizard**, in the **Recover to a Virtual Machine** drop-down menu, select **ESXi**.
 3. Enter the information for accessing the virtual machine as described in the following table, and then click **Next**.

Text Box	Description
Host name	Enter a name for the host machine.
Port	Enter the port for the host machine. The default is 443.
User name	Enter the logon credentials for the host machine.
Password	Enter the logon credentials for the host machine.

4. On the **Virtual Machine Options** page, enter the information described in the following table.

Text Box	Description
Resource Pool	Select a resource pool from the drop-down list.
Data Store	Select a data store from the drop-down list.
Virtual Machine Name	Enter a name for the virtual machine.
Memory	Click Use a specific amount of RAM to specify how much RAM to use. For example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machines. (Recommended)
Disk Provisioning	Select the type of disk provisioning as either Thin or Thick.
Disk Mapping	Specify the type of disk mapping as either Automatic or Manual.
Version	Select the version of the virtual machine.

5. Click **Next**.
6. On the **Volumes** page, select the volumes you want to export, and then click **Next**.
7. On the **Summary** page, click **Finish** to complete the wizard and start the export.



NOTE: You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Exporting Windows data using VMware workstation export

In AppAssure, you can choose to export data using VMware Workstation Export by performing a onetime or continuous export. Complete the steps in the following procedures to export using VMware Workstation Export for the appropriate type of export.

Performing a one-time VMware Workstation export



To perform a one-time VMware Workstation export:

1. In the Core Console, navigate to the machine you want to export.
2. On the **Summary** click **Actions** → **Export** → **One-time**.
The **Export Wizard** displays on the **Protected Machines** page.
3. Select a machine for export, and then click **Next**.
4. On the **Recovery Points** page, select the recovery point that you want to export, and then click **Next**.

Defining one-time settings for performing a VMware Workstation export

To define one-time settings for performing a VMware Workstation export:

1. On the **Destination** page in the **Export Wizard**, in the **Recover to Virtual machine** drop-down menu, select **VMware Workstation**, and then click **Next**.
2. On the **Virtual Machine Options** page, enter the parameters for accessing the virtual machine as described in the following table.

Text Box	Description
Location	<p>Specify the path of the local folder or network share on which to create the virtual machine.</p> <p> NOTE: If you specified a network share path, you will need to enter a valid logon credentials for an account that is registered on the target machine. The account must have read and write permissions to the network share.</p>
User Name	<p>Enter the logon credentials for the virtual machine.</p> <ul style="list-style-type: none">• If you specified a network share path, you must enter a valid user name for an account that is registered on the target machine.• If you entered a local path, a user name is not required.
Password	<p>Enter the logon credentials for the virtual machine.</p> <ul style="list-style-type: none">• If you specified a network share path, you must enter a valid password for an account that is registered on the target machine.• If you entered a local path, a password is not required.
Virtual Machine Name	<p>Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4.</p> <p> NOTE: The default name is the name of the source machine.</p>
Version	<p>Specify the version of VMware Workstation for the virtual machine. You can choose from:</p> <ul style="list-style-type: none">• VMware Workstation 7.0• VMware Workstation 8.0• VMware Workstation 9.0

Text Box	Description
Memory	<p>Specify the memory usage for the virtual machine by clicking one of the following:</p> <ul style="list-style-type: none"> • Use the same amount of RAM as the source machine - To specify that the RAM configuration is the same as the source machine. • Use a specific amount of RAM - To specify how much RAM to use; for example, 4096 Megabytes (MB). The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine. (recommended)

3. Click **Next**.
4. On the **Summary** page, click **Finish** to complete the wizard and start the export.





NOTE: You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Performing a continuous (virtual standby) VMware workstation export


To perform a continuous (virtual standby) VMware Workstation export:

1. In the Core Console, perform one of the following:
 - On the Virtual Standby tab, click **Add** to launch the **Export Wizard**. On the **Protected Machines** page of the **Export Wizard**, select the protected machine you want to export, and then click **Next**.
 - Navigate to the machine you want to export, and, on the **Summary** tab in the **Actions** drop-down menu for that machine, click **Export** → **Virtual Standby**.
2. On the **Destination** page of the **Export Wizard**, click **Recover to a Virtual Machine** → **VMware Workstation**.
3. Click **Next**.
4. On the **Virtual Machine Options** page, enter the parameters for accessing the virtual machine as described in the following table.

Text Box	Description
Target Path	<p>Specify the path of the local folder or network share on which to create the virtual machine.</p> <p> NOTE: If you specified a network share path, enter a valid logon credentials for an account that is registered on the target machine. The account must have read and write permissions to the network share.</p>
User Name	<p>Enter the logon credentials for the virtual machine.</p> <ul style="list-style-type: none"> • If you specified a network share path, you must enter a valid user name for an account that is registered on the target machine. • If you entered a local path, a user name is not required.
Password	<p>Enter the logon credentials for the virtual machine.</p> <ul style="list-style-type: none"> • If you specified a network share path, you must enter a valid password for an account that is registered on the target machine. • If you entered a local path, a password is not required.

Text Box	Description
Virtual Machine	Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4.  NOTE: The default name is the name of the source machine.
Version	Specify the version of VMware Workstation for the virtual machine. You can choose from: <ul style="list-style-type: none"> • VMware Workstation 7.0 • VMware Workstation 8.0 • VMware Workstation 9.0
Memory	Specify the memory for the virtual machine by clicking one of the following: <ul style="list-style-type: none"> • Use the same amount of RAM as the source machine - To specify that the RAM configuration is the same as the source machine. • Use a specific amount of RAM - To specify how much RAM to use; for example, 4096 Megabytes (MB). The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine.

5. Select **Perform initial ad-hoc export** to perform the virtual export immediately instead of after the next scheduled snapshot.
6. Click **Next**.
7. On the **Volumes** page, select the volumes to export, for example, C:\ and D:\, click **Next**.
8. On the **Summary** page, click **Finish** to complete the wizard and to start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Exporting Windows data using Hyper-V export

In AppAssure, you can choose to export data using Hyper-V Export by performing a one-time or continuous export. Complete the steps in the following procedures to export using Hyper-V Export for the appropriate type of export.

Performing a one-time Hyper-V export

To perform a one-time Hyper-V export:

1. In the Core Console, navigate to the machine you want to export.
2. On the Summary tab, click **Actions** → **Export** → **One-time**.
The **Export Wizard** displays on the **Protected Machines** page.
3. Select a machine for export, and then click **Next**.
4. On the **Recovery Points** page, select the recovery point that you want to export, and then click **Next**.


Defining one-time settings for performing a Hyper-V export

To define one-time settings for performing a Hyper-V export:


1. From the Hyper-V dialog box, click **Use local machine** to perform the Hyper-V export to a local machine with the Hyper-V role assigned.
2. Click the **Remote host** option to indicate that the Hyper-V server is located on a remote machine. If you selected the Remote host option, enter the parameters for the remote host described as follows:

Text Box	Description
Host Name	Enter an IP address or host name for the Hyper-V server. It represents the IP address or host name of the remote Hyper-V server.
Port	Enter a port number for the machine. It represents the port through which the Core communicates with this machine.
User Name	Enter the user name for the user with administrative privileges for the workstation with the Hyper-V server. It is used to specify the logon credentials for the virtual machine.
Password	Enter the password for the user account with administrative privileges on the workstation with Hyper-V server. It is used to specify the logon credentials for the virtual machine.

3. Click **Next**.
4. On the **Virtual Machines Options** page in the **VM Machine Location** text box, enter the path or location for the virtual machine. For example, **D:\export**. The VM location must have sufficient space to hold the VM metadata and virtual drives needed for the virtual machine.
5. Enter the name for the virtual machine in the **Virtual Machine Name** text box.
The name that you enter appears in the list of virtual machines in the Hyper-V Manager console.
6. Click one of the following:
 - **Use the same amount of RAM** as the source machine to identify that the RAM use is identical between the virtual and source machines.
 - **Use a specific amount of RAM** to specify how much memory the virtual machine has after the export; for example, 4096 MB. (recommended)
7. To specify the disk format, next to **Disk Format**, click one of the following:
 - **VHDX**
 - **VHD**

 **NOTE:** Hyper-V Export supports VHDX disk formats if the target machine is running Windows 8 (Windows Server 2012) or higher. If the VHDX is not supported for your environment, the option is disabled.
8. On the **Volumes** page, select the volume(s) to export. For the virtual machine to be an effective backup of the protected machine include the protected machine's boot drive. Example: C:\.
Your selected volumes should be no larger than 2040 GB for VHD. If the selected volumes are larger than 2040 GB, and the VHD format is selected, you will receive an error.
9. On the **Summary** page, click **Finish** to complete the wizard and to start the export.

Performing a continuous (virtual standby) Hyper-V export

-  **NOTE:** Only the 3 TB with 2 VMs configuration of DL1000 supports the one-time export and continuous export (virtual standby) capabilities.

To perform a continuous (virtual standby) Hyper-V export:

1. In the Core Console, on the **Virtual Standby** tab, click **Add** to launch the **Export Wizard**. On the **Protected Machines** page of the **Export Wizard**.
2. Select the machine you want to export and then click **Next**.
3. On the **Summary** tab, click **Export** → **Virtual Standby**.
4. From the Hyper-V dialog box, click **Use local machine** to perform the Hyper-V export to a local machine with the Hyper-V role assigned.
5. Click the **Remote host** option to indicate that the Hyper-V server is located on a remote machine. If you selected the Remote host option, enter the parameters for the remote host described as follows:

Text Box	Description
Host Name	Enter an IP address or host name for the Hyper-V server. It represents the IP address or host name of the remote Hyper-V server.
Port	Enter a port number for the machine. It represents the port through which the Core communicates with this machine.
User Name	Enter the user name for the user with administrative privileges for the workstation with the Hyper-V server. It is used to specify the logon credentials for the virtual machine.
Password	Enter the password for the user account with administrative privileges on the workstation with Hyper-V server. It is used to specify the logon credentials for the virtual machine.


6. On the **Virtual Machines Options** page in the **VM Machine Location** text box, enter the path or location for the virtual machine. For example, D:\export. The VM location must have sufficient space to hold the VM metadata and virtual drives needed for the virtual machine.
7. Enter the name for the virtual machine in the **Virtual Machine Name** text box.
The name that you enter appears in the list of virtual machines in the Hyper-V Manager console.
8. Click one of the following:
 - **Use the same amount of RAM** as the source machine to identify that the RAM use is identical between the virtual and source machines.
 - **Use a specific amount of RAM** to specify how much memory the virtual machine has after the export; for example, 4096 MB (recommended).
9. To specify the Generation, click one of the following:
 - Generation 1 (recommended)
 - Generation 2
10. To specify the disk format, next to **Disk Format**, click one of the following:
 - **VHDX** (Default)
 - **VHD**



NOTE: Hyper-V Export supports VHDX disk formats if the target machine is running Windows 8 (Windows Server 2012) or higher. If the VHDX is not supported for your environment, the option is disabled. On the Network Adapters page, select the virtual adapter to be connected to a switch.

11. On the **Volumes** page, select the volume(s) to export. For the virtual machine to be an effective backup of the protected machine include the protected machine's boot drive. Example, C:\. Your selected volumes should be no larger than 2040 GB for VHD. If the selected volumes are larger than 2040 GB, and the VHD format is selected, you will receive an error.


12. On the **Summary** page, click **Finish** to complete the wizard and to start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab

Exporting Windows data using Oracle VirtualBox export

In AppAssure, you can choose to export data using VirtualBox Export by performing a one-time or continuous export, or by establishing a continuous export (for virtual standby).

Complete the steps in the following procedures for the appropriate type of export.

 **NOTE:** To perform this type of export, you should have Oracle VirtualBox installed on the Core machine. VirtualBox Version 4.2.18 or higher is supported for Windows hosts.

Performing a one-time Oracle VirtualBox export

To perform a one-time Oracle VirtualBox export:

1. In the Core Console, navigate to the Linux machine you want to export.
2. On the **Summary** tab, click **Actions** → **Export** → **One-time**.
The **Export Wizard** displays on the **Protected Machines** page.
3. Select a machine for export, and then click **Next**.
4. On the **Recovery Points** page, select the recovery point that you want to export, and then click **Next**.
5. On the **Destination** page in the **Export Wizard**, in the **Recover to Virtual machine** drop-down menu, select **VirtualBox**, click **Next**.
6. On the **Virtual Machine Options** page, select **Remote Linux Machine**.
7. Enter the parameters for accessing the virtual machine as follows:

Text Box	Description
VirtualBox Host Name	Enter an IP address or host name for the VirtualBox server. This field represents the IP address or host name of the remote VirtualBox server.
Port	Enter a port number for the machine. This number represents the port through which the Core communicates with this machine.
Virtual Machine Name	Specify a target path to create the virtual machine.
User Name	User name of the account on the target machine, for example, root.
Password	Enter the logon credentials for the host machine.
Memory	Specify the memory for the virtual machine.

8. On the **Volumes** page, select the volumes of data to export, and then click **Next**.
9. On the **Summary** page, click **Finish** to complete the wizard and to start the export.



 **NOTE:** You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Performing a continuous (virtual standby) Oracle VirtualBox export


To perform a continuous (virtual standby) VirtualBox export:

1. In the Core Console, do one of the following:

- On the **Virtual Standby** tab, click **Add** to launch the **Export Wizard**. On the **Protected Machines** page of the **Export Wizard**, select the protected machine you want to export, and then click **Next**.
 - Navigate to the machine you want to export, and, on the **Summary** tab in the **Actions** drop-down menu for that machine, click **Export** → **Virtual Standby**.
2. On the **Destination** page in the **Export Wizard**, in the **Recover to Virtual machine** drop-down menu, select **VirtualBox**, and then click **Next**.
 3. On the **Virtual Machine Options** page, select **Use Windows machine**.
 4. Enter the parameters for accessing the virtual machine as described in the following table.

Text Box	Description
Virtual Machine Name	<p>Enter a name for the virtual machine being created.</p> <p> NOTE: The default name is the name of the source machine.</p>
Target Path	<p>Specify a local or remote target path to create the virtual machine.</p> <p> NOTE: The target path should not be a root directory.</p> <p>If you specify a network share path, you will need to enter valid logon credentials (user name and password) for an account that is registered on the target machine. The account must have read and write permissions to the network share.</p>
Memory	<p>Specify the memory for the virtual machine.</p> <ul style="list-style-type: none"> • Click Use the same amount of RAM as the source machine to specify that the RAM configuration is the same as the source machine. • Click Use a specific amount of RAM to specify how much RAM to use; for example, 4096 Megabytes (MB). The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine.

5. To specify a user account for the virtual machine, select **Specify the user account for the exported virtual machine**, and then enter the following information. This refers to a specific user account for which the virtual machine will be registered in the event there are multiple user accounts on the virtual machine. When this user account is logged on, only this user will see this Virtual Machine in VirtualBox manager. If an account is not specified, then the Virtual Machine will be registered for all existing users on the Windows machine with VirtualBox.
 - User name - Enter the user name for which the virtual machine is registered.
 - Password - Enter the password for this user account.
6. Select **Perform initial ad-hoc export** to perform the virtual export immediately instead of after the next scheduled snapshot.
7. Click **Next**.
8. On the **Volumes** page, select the volumes to export, for example, C:\ and D:\, and then click **Next**.
9. On the **Summary** page, click **Finish** to complete the wizard and to start the export.

 **NOTE:** You can monitor the status and progress of the export by viewing the **Virtual Standby** or **Events** tab.

Restoring Volumes from a Recovery Point

You can restore the volumes on a protected machine from the recovery points stored in the AppAssure Core. To restore volumes from a recovery point:

1. In the Core Console, click the **Restore** tab.

The **Restore Machine Wizard** appears.

2. From the **Protected Machines** page, select the protected machine for which you want to restore data, and then click **Next**.



NOTE: The protected machine must have the Agent software installed and must have recovery points from which you will perform the restore operation.

The **Recovery Points** page appears.

3. From the list of recovery points, search for the snapshot you want to restore to the agent machine.



NOTE: If required, use the navigation buttons at the bottom of the page to display additional recovery points. Or if you want to limit the amount of recovery points showing in the Recovery Points page of the wizard, you can filter by volumes (if defined) or by creation date of the recovery point.

4. Click any recovery point to select it, and then click **Next**.

The **Destination** page appears.

5. On the **Destination** page, choose the machine to which you want to restore data as follows:

- If you want to restore data from the selected recovery point to the same agent machine (for example, Machine1), and if the volumes you want to restore do not include the system volume, then select **Recover to a protected machine (only non-system volumes)**, verify that the destination machine (Machine1) is selected, and then click **Next**. The **Volume Mapping** page appears. Proceed to Step 7.
 - If you want to restore data from the selected recovery point to a different protected machine (for example, to replace the contents of Machine2 with data from Machine1), then select **Recover to a protected machine (only non-system volumes)**, select the destination machine (for example, Machine2) from the list, and then click **Next**. The **Volume Mapping** page appears. Proceed to Step 7.
 - If you want to restore from the selected recovery point to the same machine or a different machine using a boot CD and if the volumes you want to restore do not include the system volume, then select **Recover to any target machine using a boot CD**.
 - To continue and create the boot CD with information from the selected recovery point, click **Next** and proceed to Step 10.
 - If you have already created the boot CD and the target machine has been started using the boot CD, then proceed to Step 17.
 - If you want to restore from a recovery point to a system volume (for example, the C drive of the agent machine named Machine1), you must perform a BMR. For more information on performing a BMR for Windows, see [Launching Bare Metal Restore For Windows Machines](#).
 - For more information on performing a BMR for Linux, see Roadmap for Performing a Bare Metal Restore for Linux Machines [Launching A Bare Metal Restore For A Linux Machine](#).
6. To connect to the Universal Recovery Console (URC) on the target machine, do the following:
 - a. Select **I already have a boot CD running on the target machine**.
 - b. In the IP address text box, enter the IP address of the target machine with the boot CD.
 - c. In the Authentication Key text box, enter the authentication key from the URC on the target machine, and then click **Next**.

The **Disk Mapping** page appears. Proceed to Step 20.

7. On the **Volume Mapping** page, for each volume in the recovery point that you want to restore, select the appropriate destination volume. If you do not want to restore a volume, in the Destination Volumes column, select **Do not restore**.
8. Select **Show advanced options** and then do the following:
 - For restoring to Windows machines, if you want to use Live Recovery, select **Live Recovery**.
Using the Live Recovery instant recovery technology in AppAssure, you can instantly recover or restore data to your physical machines or to virtual machines from stored recovery points of Windows machines, which includes Microsoft Windows Storage Spaces. Live Recovery is not available for Linux machines.
 - If you want to force dismount, select **Force Dismount**.
If you do not force a dismount before restoring data, the restore may fail with a volume in use error.
9. Proceed to Step 20.
10. On the Boot CD page, do the following:
 - a. In the **Output path** text field, type the path where the boot CD ISO image should be stored,
 - b. Under **Environment** select the architecture best suited for the hardware you are restoring:
 - To restore on any Windows machine with a 64-bit architecture, select **Windows 8 64-bit**.
 - To restore on any machine with a 32-bit (x86) architecture, select **Windows 7 32-bit**.
11. Optionally, to set up network parameters for the restored agent, or to use UltraVNC, select **Show advanced options** and do one of the following:
 - To establish a network connection for the restored machine, select **Use the following IP address** as described in the following table.

Option	Description
--------	-------------

IP Address	Specify an IP address or host name for the restored machine.
Subnet Mask	Specify the subnet mask for the restored machine.
Default Gateway	Specify the default gateway for the restored machine.
DNS Server	Specify the domain name server for the restored machine.

- To define UltraVNC information, select **Add UltraVNC** as described in the following table. Use this option if you require remote access to the recovery console. You cannot log on using Microsoft Terminal Services while using the boot CD.

Option	Description
--------	-------------

Password	Specify a password for this UltraVNC connection.
Port	Specify a port for this UltraVNC connection. The default port is 5900.

12. Click **Next**.
13. To inject a driver, do the following:
 - a. Select **Add an archive of drivers**.
 - b. Navigate to a ZIP file containing the archive, select the ZIP file, and click **Open**. The archive uploads and appears in the Driver Injection page.
 - c. Then click **Next**.
14. On the ISO Image page, you can see the status as the boot CD ISO image is created. When the boot CD is successful, click **Next**.
The **Connection** page appears.
15. Start the agent machine for which you want to restore data from the boot CD.

- Boot the agent machine from an ISO image, if possible.
- If not, copy the ISO image to physical media (a CD or DVD), load the disc in the agent machine, configure the machine to load from the boot CD, and restart from the boot CD.



NOTE: You may need to change the BIOS settings of the agent machine to ensure the volume that loads first is the boot CD.

The agent machine, when started from the boot CD, displays the Universal Recovery Console (URC) interface. This environment is used to restore the system drive or selected volumes directly from the AppAssure Core. Note the IP address and authentication key credentials in the URC, which refresh each time you start from the boot CD.

16. In the Core Console on the **Connection** page, enter authentication information from the URC instance of the machine you want to restore as follows:
 - a. In the IP Address text box, enter the IP address of the machine to which you are restoring from a recovery point.
 - b. In the Authentication Key text box, enter the information from the URC.
 - c. Click **Next**.

The **Disk Mapping** page appears.

17. To map volumes manually, proceed to Step 18. To map volumes automatically, perform the following:
 - a. Select **Automatic volume mapping**.
 - b. In the **Automatic volume mapping** area, select the volumes you want to restore. If you do not wish to restore a listed volume, clear the option.



NOTE: At least one volume must be selected to perform the restore.

- c. Select the destination disk for the restore.
 - d. Click **Next**, and then proceed to Step 19.
18. If you want to map volumes manually, do the following:
 - a. Select **Manual volume mapping**.
 - b. In the **Manual volume mapping** area, from the **Destination Volumes** drop-down list for each volume, select the volume you want to restore. If you do not wish to restore a listed volume, clear the option.



NOTE: At least one volume must be selected to perform the restore.

- c. Click **Finish**.



CAUTION: If you select **Finish**, all existing partitions and data on the target drive will be removed permanently, and replaced with the contents of the selected recovery point, including the operating system and all data.

The **Restore Machine Wizard** closes, and the data is restored from the selected volumes of the recovery point to the target machine. Proceed to Step 22.

19. In the **Disk Mapping Preview** page, review the parameters of the restore actions you selected. To perform the restore, click **Finish**.



CAUTION: If you select **Finish**, all existing partitions and data on the target drive will be removed permanently, and replaced with the contents of the selected recovery point, including the operating system and all data.

The **Restore Machine Wizard** closes, and the data is restored from the selected volumes of the recovery point to the target machine. Proceed to Step 22.

20. If the volumes you want to restore contain SQL or Microsoft Exchange databases, on the **Dismount Databases** page, you are prompted to dismount them. Optionally, if you want to remount these


databases after the restore is complete, select **Automatically remount all databases after the recovery point is restored**. Click **Finish**.

21. Click **OK** to confirm the status message that the restore process has started.

22. To monitor the progress of your restore action, on the Core Console, click **Events**.

Restoring volumes for a Linux machine using the Command Line

In AppAssure, you can restore volumes on your protected Linux machines using the command-line `aamount` utility. To restore volumes for a Linux machine using the command line:

 **CAUTION:** You should not attempt to restore the system or root (/) volume.

1. Run the AppAssure `aamount` utility as root, for example:

```
sudo aamount
```

2. At the AppAssure mount prompt, enter the following command to list the protected machines:

```
lm
```

3. When prompted, enter the IP address or host name of your AppAssure Core server.

4. Enter the logon credentials, that is, the username and password, for this server.

A list displays showing the machines that this AppAssure server protects. It lists the agent machines found by line item number, host/IP address, and an ID number for the machine (for example: 293cc667-44b4-48ab-91d8-44bc74252a4f).

5. Enter the following command to list the currently mounted recovery points for the specified machine:

```
lr <machine_line_item_number>
```



NOTE: You can also enter the machine ID number in this command instead of the line item number.

A list is displayed that shows the base and incremental recovery points for that machine. This list includes a line item number, date/timestamp, location of volume, size of recovery point, and an ID number for the volume that includes a sequence number at the end (for example, "293cc667-44b4-48ab-91d8-44bc74252a4f:2"), which identifies the recovery point.

6. To select a recovery point for rollback, enter the following command:

```
r [volume_recovery_point_ID_number] [path]
```

This command rolls back the volume image specified by the ID from the Core to the specified path. The path for the rollback is the path for the device file descriptor and is not the directory to which it is mounted.



NOTE: To identify the recovery point, you can also specify a line number in the command instead of the recovery point ID number. In that case, use the agent/machine line number (from the `lm` output), followed by the recovery point line number and volume letter, followed by the path, such as, `r [machine_line_item_number] [recovery_point_line_number] [volume_letter] [path]`. In this command, `[path]` is the file descriptor for the actual volume.

For example, if the `lm` output lists three agent machines, and you enter the `lr` command for number 2, and you want to roll back the 23 recovery point volume `b` to the volume that was mounted to the directory `/mnt/data`, the command is: `r2 23 b /mnt/data`.

7. When prompted to proceed, enter `y` for Yes.

after the rollback proceeds, a series of messages appear that notify you of the status.

8. Upon a successful rollback, the `aamount` utility automatically mounts and reattach the kernel module to the rolled back volume if the target was previously protected and mounted. If not, mount the rollback volume to the local disk and then verify that the files are restored.

For example, you can use the `sudo mount` command and then the `ls` command.

Launching Bare Metal Restore for Windows machines

AppAssure provides the ability to perform a Bare Metal Restore (BMR) for your Windows machines whether the hardware is similar or dissimilar. This process encompasses creating a boot CD image, burning the image to disk, booting up the target server from disk, connecting to the recovery console instance, mapping volumes, initiating the recovery, and then monitoring the process. After the bare metal restore is complete, you can continue with the task of loading the operating system and the software applications on the restored server, followed by your unique settings and configuration.

Other circumstances in which you may choose to perform a bare metal restore include hardware upgrade or server replacement.

BMR functionality is also supported for your protected Linux machines using the command-line `aamount` utility. For more information, see [Launching A Bare Metal Restore For A Linux Machine](#).

Roadmap for performing a Bare Metal Restore for a Windows machine


To perform a BMR for a Windows machine:

1. Create a boot CD.
2. Burn the image to disk.
3. Boot the target server from the boot CD.
4. Connect to the recovery disk.
5. Map the volumes.
6. Initiate the recovery.
7. Monitor the progress.

Creating a bootable CD ISO image

To perform a BMR for a Windows machine, you must create a bootable CD/ISO image in the Core Console, which contains the AppAssure Universal Recovery Console interface. The AppAssure Universal Recovery Console is an environment used to restore the system drive or the entire server directly from the AppAssure Core.

The ISO image that you create is tailored to the machine being restored; therefore, it must contain the correct network and mass storage drivers. If you anticipate that you will be restoring to different hardware from the machine on which you are creating the boot CD, you must include storage controller and other drivers in the boot CD see [Injecting Drivers in a Boot CD](#).

 **NOTE:** The International Organization for Standardization (ISO) is an international body of representatives from various national organizations who determine and set file system standards. The ISO 9660 is a file system standard that is used for optical disk media for the exchange of data. It supports various operating systems, such as Windows. An ISO image is the archive file or disk image, which contains data for every sector of the disk as well as the disk file system.

To create a bootable CD ISO image:

1. From the Core Console on which the server you want to restore is located, select the **Core** and then click the **Tools** tab.
2. Click **Boot CDs**.
3. Select **Actions**, and then click **Create Boot ISO**.


The **Create Boot CD** dialog box displays. To complete the dialog box, use the following procedures.

Naming the boot cd file and setting the path

To name the boot CD file and set the path:

In the **Create Boot CD** dialog box, enter the ISO path where to store the boot image on the Core server.


If the share on which you want to store the image is low on disk space, you can set the path as needed; for example, D:\filename.iso.

 **NOTE:** The file extension must be .iso. When specifying the path, use only alphanumeric characters, the hyphen, and the period (only to separate host names and domains). The letters a to z are case-insensitive. Do not use spaces. No other symbols or punctuation characters are permitted.

Creating connections

To create connections:


1. In **Connection Options** do one of the following:
 - To obtain the IP address dynamically using Dynamic Host Configuration Protocol (DHCP), select **Obtain IP address automatically**.
 - Optionally, to specify a static IP address for the recovery console, select **Use the following IP address** and enter the IP address, subnet mask, default gateway, and DNS server in the appropriate fields. You must specify all of these fields.
2. If required, in the **UltraVNC Options**, select **Add UltraVNC** and then enter the UltraVNC options. The UltraVNC settings enable you to manage the recovery console remotely while it is in use.

 **NOTE:** This step is optional. If you need remote access to the recovery console, you must configure and use the UltraVNC. You cannot log on using Microsoft Terminal Services while using the boot CD.

Injecting drivers in a boot cd

Driver injection is used to facilitate the operability between the recovery console, network adapter, and storage on the target server.

If you anticipate restoring to dissimilar hardware, you must inject storage controller, RAID, AHCI, chipset and other drivers in the boot CD. These drivers make it possible for the operating system to detect and operate all devices successfully.

 **NOTE:** Keep in mind that the boot CD will automatically contain Windows 7 PE 32-bit drivers.

To inject drivers in a boot CD:

1. Download the drivers from the manufacturer's website for the server and unpack them.
2. Compress the folder that contains the drivers using a file compressing utility, such as WinZip.
3. In the **Create Boot CD** dialog box, in the **Drivers** pane, click **Add a Driver**.
4. To locate the compressed driver file, navigate through the filing system. Select the file, and then click **Open**.


The injected drivers appear highlighted in the **Drivers** pane.

Creating the boot cd

To create a boot CD, after you have named the boot CD and specified the path, created a connection and optionally injected the drivers, from the **Create Boot CD** screen, click **Create Boot CD**. The ISO image is then created.

Viewing the iso image creation progress

To view the ISO image creation progress, select the **Events** tab, and then under **Tasks**, you can monitor the progress for building the ISO image.

 **NOTE:** You can also view the progress of the creation of the ISO image in the **Monitor Active Task** dialog box.

When the creation of the ISO image is complete, it is available on the **Boot CDs** page, accessible from the **Tools** menu.

Accessing the iso image

To access the ISO image, navigate to the output path you specified, or you can click the link to download the image to a location from which you can then load it on the new system. For example, network drive.

Loading a boot CD

When you have created the boot CD image, boot the target server with the newly created boot CD.

 **NOTE:** If you created the boot CD using DHCP, note the IP address and password.

To load a boot CD:

1. Navigate to the new server, load the boot CD, and then start the machine.
2. Specify to **Boot from CD-ROM**, which loads the following:
 - Windows 7 PE
 - AppAssure Agent software

The AppAssure Universal Recovery Console starts and displays the IP address and authentication password for the machine.


3. Record the IP address displayed in the Network Adapters Settings pane and the authentication password displayed in the Authentication pane. You will use this information later during the data recovery process to log back on to the console.
4. If you want to change the IP address, select it and click **Change**.

 **NOTE:** If you specified an IP address in Create Boot CD dialog box, the Universal Recovery Console uses it and displays it in the **Network Adapter settings** screen.

Injecting drivers to your target server

If you are restoring to dissimilar hardware, you must inject storage controller, RAID, AHCI, chipset and other drivers if they are not already on the boot CD. These drivers make it possible for the operating system to operate all devices on your target server successfully.

If you are unsure which drivers your target server requires, click the System Info tab in the Universal Recovery Console. This tab shows all system hardware and device types for the target server to which you want to restore.

 **NOTE:** Keep in mind that your target server automatically contains Windows 7 PE 32-bit drivers.



To inject drivers to your target server:

1. Download the drivers from the manufacturer's website for the server and unpack them.
2. Compress the folder that contains the drivers by using a file compressing utility (for example, Win Zip) and copy it to the target server.
3. In the Universal Recovery Console, click **Driver Injection**.
4. To locate the compressed driver file, navigate through the filing system and select the file.
5. If you clicked **Driver Injection** in step 3, click **Add Driver**. If you clicked **Load driver** in step 3, click **Open**.

The selected drivers are injected and will be loaded to the operating system after you reboot the target server.

Launching a restore from the Core

To launch a restore from the Core:

1. If the NICs on any system being restored are teamed (bonded), remove all but one of the network cables.
 **NOTE:** AppAssure Restore does not recognize teamed NICs. The process is not able to resolve which NIC to use if presented with more than one active connection.
2. Navigate back to the Core server and open the Core Console.
3. On the **Machines** tab, select the machine from which you want to restore data.
4. Click the **Actions** menu for the machine, click **Recovery Points** to view a list of all recovery points for that machine.
5. Expand the recovery point from which you want to restore, then click **Rollback**.
6. In the **Rollback** dialog box, under Choose **Destination**, select **Recovery Console Instance**.
7. In the **Host** and **Password** text boxes, enter the IP address and the authentication password for the new server to which you want to restore data.
 **NOTE:** The Host and Password values are the credentials you recorded in the previous task. For more information, see [Loading A Boot CD](#).
8. Click **Load Volumes** to load the target volumes to the new machine.

Mapping volumes

You can choose to map volumes to the disks on the target server automatically or manually. For automatic disk alignment, the disk is cleaned and repartitioned and all data is deleted. The alignment is performed in the order the volumes are listed and the volumes are allocated to the disks appropriately according to size, and so on. Multiple volumes can use a disk. If you manually map the drives, you cannot use the same disk twice.

For manual mapping, you must already have the new machine correctly formatted before restoring it.

To map volumes:

1. To automatically map volumes, do the following:
 - a. On the **Disk Mapping** page of the **Restore Machine Wizard**, select the **Automatically Map Volumes** tab.
 - b. In the **Disk Mapping** area, under **Source Volume**, verify that the source volume is selected and that the appropriate volumes are both listed beneath and are selected.
 - c. If the destination disk that is automatically mapped is the correct target volume, select **Destination Disk**.


- d. Click **Restore**, and then proceed to step 3.
2. To manually map volumes, do the following:
 - a. On the **Disk Mapping** page of the **Restore Machine Wizard**, select the **Manually Map Volumes** tab.
 - b. In the **Volume Mapping** area, under **Source Volume**, verify that the source volume is selected and that the appropriate volumes are both listed beneath and are selected.
 - c. Under **Destination**, from the drop-down menu, select the appropriate destination that is the target volume to perform the bare metal restore of the selected recovery point, and then click **Rollback**.
3. In the **RollbackURC** confirmation dialog box, review the mapping of the source of the recovery point and the destination volume for the rollback. To perform the rollback, click **Restore**.



CAUTION: If you select **Begin Rollback**, all existing partitions and data on the target drive will be permanently removed, and replaced with the contents of the selected recovery point, including the operating system and all data.

Viewing the recovery progress

To view the recovery progress:

1. After you initiate the rollback process, the **Active Task** dialog box displays, showing that the rollback action initiated.
 **NOTE:** This appearance of the **Active Task** dialog box does not indicate successful completion of the task.
2. Optionally, to monitor the rollback task progression, from the Active Task dialog box, click **Open Monitor Window**. You can view the status of the recovery as well as the start and end times from the **Monitor Open Task** window.



NOTE: To return to the recovery points for the source machine from the **Active Task** dialog box, click **Close**.

Starting the restored target server

To start the restored target server:

1. Navigate back to the target server, and in the **AppAssure Universal Recovery Console** interface, click **Reboot** to start the machine.
2. Specify to start Windows normally.
3. Log on to the machine.

The system is restored to its state prior to the bare metal restore.

Repairing startup problems

Keep in mind that if you restored to dissimilar hardware, you must have injected storage controller, RAID, AHCI, chipset and other drivers if they are not already on the boot CD. These drivers make it possible for the operating system to operate all devices on your target server successfully.

To repair startup problems:

1. If you encounter problems when starting the restored target server, open the Universal Recovery Console by reloading the boot CD.
2. In the Universal Recovery Console, click **Driver Injection**.
3. In the Driver Injection dialog, click **Repair Boot Problems**.
The startup parameters in the target server boot record are automatically repaired.
4. In the Universal Recovery Console, click **Reboot**.

Launching a bare metal restore for a Linux machine

Your DL1000 can perform a Bare Metal Restore (BMR) for a Linux machine including rollback of the system volume. Using the AppAssure command line utility `aamount`, roll back to the boot volume base image. Before you can perform a BMR for a Linux machine, you first must do the following:

- Obtain a BMR Live CD file from AppAssure support, which includes a bootable version of Linux.



NOTE: You can also download the Linux Live CD file from the license portal at <https://licenseportal.com>.

- Ensure that there is enough space on the hard drive to create destination partitions on the target machine to contain the source volumes. Any destination partition should be at least as large as the original source partition.
- Identify the path for the rollback, which is the path for the device file descriptor. To identify the path for the device file descriptor, use the `fdisk` command from a terminal window.



NOTE: Before you begin utilizing the AppAssure commands, you can install the screen utility. The screen utility enables you to scroll the screen to view larger amounts of data, such as a list of recovery points.

To perform a bare metal restore for a Linux machine:

1. Using the Live CD file you receive from AppAssure, boot up the Linux machine and open a Terminal window.
2. If needed, create a new disk partition, for example, by running the `fdisk` command as root, and make this partition bootable by using the `a` command.
3. Run the AppAssure `aamount` utility as root, for example:

```
sudo aamount
```

4. At the AppAssure mount prompt, enter the following command to list the protected machines:
`lm`
5. When prompted, enter the IP address or host name of your AppAssure Core server.
6. Enter the logon credentials, that is, the username and password, for this server.

A list is displayed showing the machines protected by this AppAssure Core server. It lists the machines found by line item number, host/IP address, and an ID number for the machine (for example: `293cc667-44b4-48ab-91d8-44bc74252a4f`).

7. To list the currently mounted recovery points for the machine that you want to restore, enter the following command:

```
lr <machine_line_item_number>
```



NOTE: You can also enter the machine ID number in this command instead of the line item number.

A list is displayed that shows the base and incremental recovery points for that machine. This list includes a line item number, date/timestamp, location of volume, size of recovery point, and an ID number for the volume that includes a sequence number at the end (for example: `"293cc667-44b4-48ab-91d8-44bc74252a4f:2"`), which identifies the recovery point.


8. To select the base image recovery point for rollback, enter the following command:

```
r <volume_base_image_recovery_point_ID_number> <path>
```



CAUTION: You must ensure that the system volume is not mounted.


This command rolls back the volume image specified by the ID from the Core to the specified path. The path for the rollback is the path for the device file descriptor and is not the directory to which it is mounted.


 **NOTE:** You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. Use the agent/machine line number (from the `lm` output), followed by the recovery point line number and volume letter, followed by the path, such as, `r <machine_line_item_number> <base_image_recovery_point_line_number> <volume_letter> <path>`. In this command, `<path>` is the file descriptor for the actual volume.

9. When prompted to proceed, enter `y` for Yes.

After the rollback proceeds, a series of messages appear that notify you of the status.

10. Upon a successful rollback, if needed, update the main boot record with the restored bootloader.

 **NOTE:** Repairing or setting up the bootloader is only needed if this disk is new. If this is a simple rollback to the same disk, setting up the bootloader is not necessary.

 **CAUTION:** Do not unmount a protected Linux volume manually. In the event that you need to manually unmount a protect Linux volume, you must execute the following command before unmounting the volume: `bsctl -d <path to volume>`

In this command, `<path to volume>` does not refer to the mount point of the volume but instead refers to the file descriptor of the volume; it must be in a form similar to this example: `/dev/sda1`.

Installing the screen utility

Before you begin utilizing the AppAssure commands, you can install the screen utility. The screen utility enables you to scroll the screen to view larger amounts of data, such as a list of recovery points.


To install the screen utility:

1. Using the Live CD file, start the Linux machine.
A terminal window opens.
2. Enter the following command: `sudo apt-get install screen`.
3. To start the screen utility, type `screen` at the command prompt.

Creating bootable partitions on a Linux machine

To create bootable partitions on a Linux machine by using the command line:

1. Attach to all devices using the **bsctl** utility with the following command as root: `sudo bsctl --attach-to-device /dev/<restored volume>`

 **NOTE:** Repeat this step for each restored volume.

2. Mount each restored volume by using the following commands:

```
mount /dev/<restored volume> /mnt
```

```
mount /dev/<restored volume> /mnt
```

 **NOTE:** Some system configurations may include the boot directory as part of the root volume.

3. Mount snapshot metadata for each restored volume by using the following commands:

```
sudo bsctl --reset-bitmap-store /dev/<restored volume>
```

```
sudo bsctl --map-bitmap-store /dev/<restored volume>
```

4. Verify that the Universally Unique Identifier (UUID) contains the new volumes by using either the `blkid` command or the `ll /dev/disk/by-uuid` command.
5. Verify that `/etc/fstab` contains the correct UUIDs for the root and boot volumes.
6. Install Grand Unified Bootloader (GRUB) by using the following commands:

```
mount --bind /dev/ /mnt/dev  
  
mount --bind /proc/ /mnt/proc  
  
chroot/mnt/bin/bash  
  
grub-install/dev/sda
```
7. Verify that the `/boot/grub/grub.conf` file contains the correct UUID for the root volume, or update it as needed by using a text editor.
8. Remove the Live CD disk from the CD-ROM drive and restart the Linux machine.

Replicating recovery points

Replication

Replication is the process of copying recovery points and transmitting them to a secondary location for the purpose of disaster recovery. The process requires a paired source-target relationship between two cores. Replication is managed on a per-protected-machine basis; meaning, backup snapshots of a protected machine are replicated to the target replica core. When replication is set up, the source core asynchronously and continuously transmits the incremental snapshot data to the target core. You can configure this outbound replication to your company's own data center or remote disaster recovery site (that is, a "self-managed" target core) or to a managed service provider (MSP) providing off-site backup and disaster recovery services. When you replicate to an MSP, you can use built-in workflows that let you request connections and receive automatic feedback notifications.

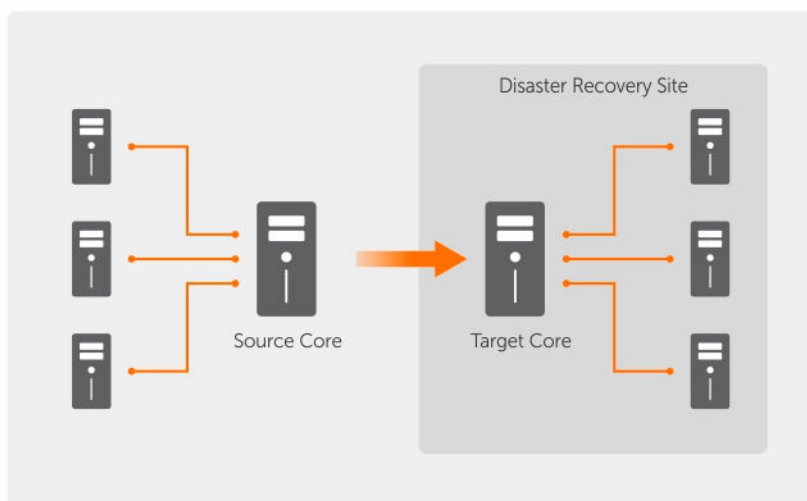


Figure 5. Basic Replication Architecture

Replication begins with seeding. The initial transfer of deduplicated base images and incremental snapshots of the protected agents, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target core by using external media. This is typically useful for large sets of data or sites with slow links. The data in the seeding archive is compressed, encrypted and deduplicated. If the total size of the archive is larger than the space available on the removable media, the archive can span across multiple devices based on the available space on the media. During the seeding process, the incremental recovery points replicate to the target site. After the target core consumes the seeding archive, the newly replicated incremental recovery points automatically synchronize.

Roadmap for performing replication


To replicate data using AppAssure, you must configure the source and target cores for replication. After you configure replication, you can then replicate data of the protected machine, monitor and manage replication, and perform recovery.

Performing replication in AppAssure involves performing the following operations:

- Configure self-managed replication. For more information on replicating to a self-managed target core, see [Replicating To A Self-Managed Core](#).
- Configure third-party replication. For more information on replicating to a third-party target core, see [Replicating To A Core Managed By A Third Party](#).
- Replicate a new protected machine attached to the source core. For more information on replicating a protected machine, see [Replicating A New Protected Machine](#).
- Replicate an existing protected machine. For more information on configuring an agent for replication, see [Replicating Agent Data On A Machine](#).
- Set replication priority for an agent. For more information on prioritizing the replication of agents, see [Setting Replication Priority For An Agent](#).
- Monitor replication as needed. For more information on monitoring replication, see [Monitoring Replication](#).
- Manage replication settings as needed. For more information on managing replication settings, see [Managing Replication Settings](#).
- Recover replicated data in the event of disaster or data loss. For more information on recovering replicated data, see [Recovering Replicated Data](#).

Replicating to a self-managed core

A self-managed core is a core to which you have access, often because it is managed by your company at an off-site location. Replication can be completed entirely on the source core, unless you choose to seed your data. Seeding requires that you consume the seed drive on the target core after you configure replication on the source core.

 **NOTE:** This configuration applies to replication to an off-site location and to mutual replication. The Core must be installed on all source and target machines. If you are configuring your system for multi-point to point replication, you must perform this task on all source cores and the one target core.

Configuring the source core to replicate to a self-managed target core

To configure the source core to replicate to a self-managed target core:

1. In the Core, click the **Replication** tab.
2. Click **Add Target Core**.
The **Replication** wizard appears.
3. Select **I have my own Target Core**, and then enter the information as described in the following table.

Text Box	Description
Host Name	Enter the host name or IP address of the Core machine to which you are replicating.

Text Box	Description
Port	Enter the port number on which the AppAssure Core communicates with the machine. The default port number is 8006.
User Name	Enter the user name for accessing the machine. For example, Administrator .
Password	Enter the password for accessing the machine.

If the Core you want to add has been paired with this source core previously, perform the following:

- a. Select **Use an existing target core**.
 - b. Select the target core from the drop-down list.
 - c. Click **Next**.
 - d. Skip to step 7.
4. Click **Next**.
 5. On the **Details** page, enter a name for this replication configuration; for example, SourceCore1. If you are re-initiating or repairing a previous replication configuration, select **My Core has been migrated and I would like to repair replication**
 6. Click **Next**.
 7. On the **Agents** page, select the agents you want to replicate, and then use the drop-down lists in the **Repository** column to select a repository for each agent.
 8. If you plan to perform the seeding process for the transfer of the base data, complete the following steps:



NOTE: Because large amounts of data need to be copied to the portable storage device, an eSATA, USB 3.0, or other high-speed connection to the portable storage device is recommended.

- a. On the **Agents** page, select **Use a seed drive to perform initial transfer**. If you currently have one or more machines replicating to a target core, you can include these protected machines on the seed drive by selecting **With already replicated**.
 - b. Click **Next**.
 - c. On the **Seed Drive Location** page, use the **Location type** drop-down list to select one of the following:
 - **Local:** In the **Location** text box, enter where you want to save the seed drive; for example, D:\work\archive.
 - **Network:** In the **Location** text box, enter where you want to save the seed drive, and then enter your credentials for the network share in the **User name** and **Password** text boxes.
 - **Cloud:** In the **Account** text box, select the account. To select a cloud account, you must first have added it in the Core Console. For more information, see [Adding A Cloud Account](#). Select the **Container** associated with your account. Select the **Folder Name** to which the archived data is to be saved.
 - d. Click **Next**.
9. In the **Seed Drive Option** dialog box, enter the information described as follows:

Text Box	Description
Maximum Size	<p>Large archives of data can be divided into multiple segments. Select the maximum size of the segment you want to reserve for creating the seed drive by doing one of the following:</p> <ul style="list-style-type: none"> • Select Entire Target to reserve all available space in the path provided on the Seed Drive Location page for future use (for example, if the location is D:\work\archive, all of the available space on the D: drive is reserved if

Text Box	<p>Description</p> <p>required for copying the seed drive, but is not reserved immediately after starting the copying process).</p> <ul style="list-style-type: none"> • Select the blank text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve.
Customer ID (optional)	Optionally, enter the customer ID that was assigned to you by the service provider.
Recycle action	<p>In the event that the path already contains a seed drive, select one of the following options:</p> <ul style="list-style-type: none"> • Do not reuse — Does not overwrite or clear any existing data from the location. If the location is not empty, the seed drive write fails. • Replace this core — Overwrites any pre-existing data pertaining to this core but leave the data for other cores intact. • Erase completely — Clears all data from the directory before writing the seed drive.
Comment	Enter a comment or description of the archive.
Add all Agents to Seed Drive	Select the agents you want to replicate using the seed drive.
Build RP chains (fix orphans)	<p>Select this option to replicate the entire recovery point chain to the seed drive. This option is selected by default.</p> <p>Typical seeding in AppAssure replicates only the latest recovery point to the seed drive, which reduces the amount of time and space required for creating the seed drive. Opting to build recovery point (RP) chains to the seed drive requires enough space on the seed drive to store the latest recovery points from the specified agent or agents, and may take additional time to complete the task.</p>
Use compatible format	Select this option to create the seed drive in a format that is compatible with both new and older versions of the AppAssure Core.

10. On the **Agents** page, select the agents you want to replicate to the target core using the seed drive.

11. Click **Finish**.

12. If you created a seed drive, send it to your target core.

The pairing of the source core to the target core is complete. Replication begins, but produces orphaned recovery points on the target core until the seed drive is consumed and provides the necessary base images.

Consuming the seed drive on a target core

This procedure is only necessary if you created a seed drive while Configuring Replication For A Self-Managed Core.

To consume the seed drive on a target core:

- 1.** If the seed drive was saved to a portable storage device such as a USB drive, connect the drive to the target core.
- 2.** From the Core Console on the target core, select the **Replication** tab.

3. Under **Incoming Replication**, select the correct source core by using the drop-down menu, and then click **Consume**.

The Consume window appears.

4. For **Location type**, select one of the following options from the drop-down list:

- Local
- Network
- Cloud

5. Enter the following information as needed:

Text Box	Description
Location	Enter a path to where the seed drive is located, such as a USB drive or a network share (for example, D:\).
User name	Enter the user name for the shared drive or folder. User name is required only for a network path.
Password	Enter the password for the shared drive or folder. Password is required only for a network path.
Account	Select an account from the drop-down list. To select a cloud account, you must first have added it in the Core Console.
Container	Select a container associated with your account from the drop-down menu.
Folder Name	Enter the name of the folder in which the archived data is saved; for example, - Archive-[DATE CREATED]- [TIME CREATED]

6. Click **Check File**.

After the Core checks the file, it automatically populates the **Date Range** with the dates of the oldest and newest recovery points contained in the seed drive. It also imports any comments entered in Configuring Replication For A Self-Managed Core.

7. Under **Agent Names** on the **Consume** window, select the machines for which you want to consume data, and then click **Consume**.



NOTE: To monitor the data consumption progress, select the **Events** tab.

Abandoning an outstanding seed drive

If you create a seed drive with the intent to consume it on the target core but choose not to send it to the remote location, a link for the outstanding seed drive remains on the source core **Replication** tab. You may want to abandon the outstanding seed drive in favor of different or more current seed data.



NOTE: This procedure removes the link to the outstanding seed drive from the Core Console on the source core. It does not remove the drive from the storage location on which it is saved.

To abandon an outstanding seed drive:

1. From the Core Console on the source core, select the **Replication** tab.
2. Click **Outstanding Seed Drive (#)**.

The **Outstanding seed drives** section appears. It includes the name of the remote target core, the data and time at which the seed drive was created, and the data range of the recovery points included on the seed drive.


3. Click the drop-down menu for the drive that you want to abandon, then select **Abandon**. The **Outstanding Seed Drive** window appears.

4. Click **Yes** to confirm the action.

The seed drive is removed. If there are no more seed drives that exist on the source core, then the next time that you open the **Replication** tab, the **Outstanding Seed Drive (#)** link and **Outstanding seed drives** section do not appear.

Replicating to a core managed by a third party

A third-party core is a target core that is managed and maintained by an MSP. Replicating to a core managed by a third party does not require you to have access to the target core. After a customer configures replication on the source core or cores, the MSP completes the configuration on the target core.

 **NOTE:** This configuration applies to hosted and cloud replication. The AppAssure Core must be installed on all source core machines.

Replicating a new agent

When you add an AppAssure Agent for protection on a source core, AppAssure gives you the option to replicate the new agent to an existing target core.

To replicate a new agent:

1. Navigate to the Core Console, and then click the **Machines** tab.
2. In the **Actions** drop-down menu, click **Protect Machine**.
3. In the **Protect Machine** dialog box, enter the information as described in the following table.


Text Box	Description
Host	Enter the host name or IP address of the machine that you want to protect.
Port	Enter the port number the AppAssure Core uses to communicate with the agent on the machine.
Username	Enter the username used to connect to this machine. For example, Administrator.
Password	Enter the password used to connect to this machine.

4. Click **Connect** to connect to this machine.
5. Click **Show Advanced Options**, and edit the following settings as needed.

Text Box	Description
Display Name	Enter a name for the machine to be displayed in the Core Console.
Repository	Select the repository on the AppAssure Core where the data from this machine is stored.
Encryption Key	Specify whether encryption is applied to the data for every volume on this machine stored in the repository.

 **NOTE:** The encryption settings for a repository are defined under the **Configuration** tab in the Core Console.

Remote Core	Specify the target core to which you want to replicate the agent.
--------------------	---

Text Box	Description
Remote Repository	The name of the desired repository on the target core in which to store the replicated data from this machine.
Pause	Select this check box if you want to pause replication; for example, to pause it until after AppAssure takes a base image of the new agent.
Schedule	Select one of the following options: <ul style="list-style-type: none"> • Protect all volumes with default schedule • Protect specific volumes with custom schedule <p> NOTE: The default schedule is every 15 minutes.</p>
Initially pause protection	Select this check box if you want to pause protection; for example, to prevent AppAssure from taking the base image until after peak usage hours.

6. Click **Protect**.

Replicating agent data on a machine

Replication is the relationship between the target and source cores in the same site, or across two sites with slow link on a per agent basis. When replication is set up between two cores, the source core asynchronously transmits the incremental snapshot data of select agents to the target or source core. Outbound replication can be configured to a Managed Service Provider providing off-site backup and disaster recovery service or to a self-managed core. To replicate agent data on a machine:

1. From the Core Console, click the **Machines** tab.
2. Select the machine that you want to replicate.
3. In the **Actions** drop-down menu, click **Replication**, and then complete one of the following options:
 - If you are setting up replication, click **Enable**.
 - If you already have an existing Replication set up, click **Copy**.

The **Enable Replications** dialog box appears.
4. In the **Host** text box, enter a host name.
5. Under **Agents**, select the machine that has the agent and data that you want to replicate.
6. If needed, select the check box **Use a seed drive to perform initial transfer**.
7. Click **Add**.
8. To pause or resume the replication, click **Replication** in the **Actions** drop-down menu, and then click **Pause** or **Resume** as needed.

Setting replication priority for an agent

To set replication priority for an agent:

1. From the Core Console, select the protected machine for which you want to set replication priority, and click the **Configuration** tab.
2. Click **Select Transfer Settings**, and then use the **Priority** drop-down list to select one of the following options:
 - **Default**
 - **Highest**
 - **Lowest**

- 1
- 2
- 3
- 4



NOTE: The default priority is 5. If one agent is given the priority 1, and another agent is given the priority Highest, the agent with the Highest priority replicates before the agent with the 1 priority.

3. Click **OK**.

Monitoring replication

When replication is set up, you can monitor the status of replication tasks for the source and target cores. You can refresh status information, view replication details, and more.

To monitor replication:

1. In the Core Console, click the **Replication** tab.
2. On this tab, you can view information about and monitor the status of replication tasks described as follows:

Table 4. Monitoring replication

Section	Description	Available Actions
Pending Replication Requests	Lists your customer ID, email address, and host name when a replication request is submitted to a third-party service provider. It is listed here until the MSP accepts the request.	In the drop-down menu, click Ignore to ignore or reject the request.
Outstanding Seed Drives	Lists seed drives that have been written but not yet consumed by the target core. It includes the remote core name, date on which it was created, and the date range.	In the drop-down menu, click Abandon to abandon or cancel the seed process.
Outgoing Replication	Lists all target cores to which the source core is replicating. It includes the remote core name, the state of existence, the number of protected machines being replicated, and the progress of a replication transmission.	On a source core, in the drop-down menu, you can select the following options: <ul style="list-style-type: none"> • Details — Lists the ID, URI, display name, state, customer ID, email address, and comments for the replicated core. • Change Settings — Lists the display name and lets you edit the host and port for the target core. • Add Agents — Lets you select a host from a drop-down list, select protected machines for replication,

Section	Description	Available Actions
Incoming Replication	Lists all source machines from which the target receives replicated data. It includes the remote core name, state, machines, and progress.	<p>and create a seed drive for the new protected machine's initial transfer.</p> <p>On a target core, in the drop-down menu, you can select the following options:</p> <ul style="list-style-type: none"> • Details — Lists the ID, host name, customer ID, email address, and comments for the replicated core. • Consume — Consumes the initial data from the seed drive and saves it to the local repository.

3. Click the **Refresh** button to update the sections of this tab with the latest information.

Managing replication settings

You can adjust a number of settings for how replication executes on the source and target cores. To manage replication settings:

1. In the Core Console, click the **Replication** tab.
2. In the **Actions** drop-down menu, click **Settings**.
3. In the **Replication Settings** window, edit the replication settings described as follows:

Option	Description
Cache lifetime	Specify the amount of time between each target-core status request performed by the source core.
Volume image session timeout	Specify the amount of time the source core spends attempting to transfer a volume image to the target core.
Max. concurrent replication jobs	Specify the number of protected machines permitted to replicate to the target core at one time.
Max. parallel streams	Specify the number of network connections permitted to be used by a single protected machine to replicate that machine's data at one time.

4. Click **Save**.

Removing replication

You can discontinue replication and remove protected machines from replication in several ways. The options include:

- [Removing An Agent From Replication On The Source Core](#)
- [Removing An Agent On The Target Core](#)
- [Removing A Target Core From Replication](#)
- [Removing A Source Core From Replication](#)



NOTE: Removing a source core results in the removal of all replicated machines that are protected by that core.

Removing a protected machine from replication on the source Core

To remove a protected machine from replication on the source core:

1. From the source core, open the Core Console, and click the **Replication** tab.
2. Expand the **Outgoing Replication** section.
3. In the drop-down menu for the protected machine that you want to remove from replication, click **Delete**.
4. In the **Outgoing Replication** dialog box, click **Yes** to confirm deletion.

Removing a protected machine on the target Core

To remove a protected machine on the target core:

1. On the target core, open the Core Console, and click the **Replication** tab.
2. Expand the **Incoming Replication** section.
3. In the drop-down menu for the protected machine that you want to remove from replication, click **Delete**, and then select one of the following options.

Option	Description
Relationship Only	Removes the protected machine from replication but retains the replicated recovery points.
With Recovery Point	Removes the protected machine from replication and deletes all replicated recovery points received from that machine.

Removing a target Core from replication

To remove a target core from replication:

1. On the source core, open the Core Console, and click to the **Replication** tab.
2. Under **Outgoing Replication**, click the drop-down menu next to the remote core that you want to delete, and click **Delete**.
3. In the **Outgoing Replication** dialog box, click **Yes** to confirm deletion.

Removing a source Core from replication



NOTE: Removing a source core results in the removal of all replicated agents protected by that core.

To remove a source core from replication:

1. On the target core, open the Core Console, and click the **Replication** tab.
2. Under **Incoming Replication**, in the drop-down menu, click **Delete**, and then select one of the following options.

Option	Description
Relationship Only	Removes the source core from replication but retains the replicated recovery points.
With Recovery Points	Removes the source core from replication and deletes all replicated recovery points received from that machine.

3. In the **Incoming Replication** dialog box, click **Yes** to confirm deletion.

Recovering replicated data

Day-to-day replication functionality is maintained on the source core, while only the target core is capable of completing the functions necessary for disaster recovery.

For disaster recovery, the target core can use the replicated recovery points to recover the protected agents and core.

You can perform the following recovery options from the target core:

- Mount recovery points.
- Roll back to recovery points.
- Perform a virtual machine (VM) export.
- Perform a bare metal restore (BMR).
- Perform Failback (in the event you have a Failover/Failback replication environment set up).

Understanding failover and failback

AppAssure supports failover and failback in replicated environments, in case of a severe outage in which the source core and agents fail. Failover refers to switching to a redundant or standby target (AppAssure Core) upon system failure or abnormal termination of a source core and associated agents. The main goal of failover is to launch a new agent identical to the failed agent. The secondary goal is to switch the target core into a new mode so that the target core protects the failover agent in the same way as the source core protected the initial agent before the failure. The target core can recover instances from replicated agents and immediately commence protection on the failed-over machines.

Failback is the process of restoring an agent and core back to their original states (before failure). The primary goal of failback is to restore the agent (in most cases, this is a new machine replacing a failed agent) to a state identical to the latest state of the new, temporary agent. When restored, it is protected by a restored source core. Replication is also restored, and the target core acts as a replication target again.

Performing Failover

When you encounter a disaster situation in which your source core and associated agents have failed, you can enable failover in AppAssure to switch protection to your identical failover (target) core. The target core becomes the only core protecting the data in your environment, and you then launch a new agent to temporarily replace the failed agent.

To perform failover on the target core:

1. Navigate to the Core Console on the target core, and click the **Replication** tab.
2. Under **Incoming Replication**, select the source core, and then expand the details under the individual agent.


3. On the **Actions** menu for that core, click **Failover**.
The **Fail Over** dialog box appears and lists the next steps required for completing a failover.
4. Click **Continue**.
5. In the left navigation area, under **Protected Machines**, select the machine that has the associated AppAssure Agent software with recovery points.
6. Export the backup recovery point information on that agent to a virtual machine.
7. Export the backup recovery point information on that agent to a virtual machine.
8. Start the virtual machine that now includes the exported backup information.
You need to wait for the device driver software to be installed.
9. Reboot the virtual machine and wait for the agent service to start.
10. Go back to the Core Console for the target core and verify that the new agent is displayed under **Protected Machines** and on the **Replication** tab under **Incoming Replication**.
11. Force multiple snapshots, and verify they complete correctly.
For more information, see [Forcing a Snapshot](#).
12. You can now proceed with performing failback.
For more information, see [Performing Failback](#).

Performing Failback

After you repair or replace the failed original source core and agents, you need to move the data from your failed-over machines to restore the source machines.

To perform failback:

1. Navigate to the Core Console on the target core, and click the **Replication** tab.
2. Under **Incoming Replication**, select the failover agent and expand the details.
3. On the **Actions** menu, click **Failback**.
The **Fail Back** dialog box opens to describe the steps you need to follow before you click the **Continue** button to complete failback.
4. Click **Cancel**.
5. If the failed-over machine is running Microsoft SQL Server or Microsoft Exchange Server, stop those services.
6. Force a snapshot of the machine. For more information, see [Forcing a Snapshot](#).
7. Shut down the failed-over machine.
8. Create an archive of the failed-over agent and output it to disk or a network share location.
For more information about creating archives, see [Creating An Archive](#).
9. After you create the archive, navigate to the Core Console on the newly repaired source core, and click the **Tools** tab.
10. Import the archive you just created in Step 8.
For more information, see [Importing An Archive](#).
11. Go back to the Core Console on the target core, and click the **Replication** tab.
12. Under **Incoming Replication**, select the failover agent and expand the details.
13. In the **Failback** dialog box, click **Continue**.
14. Shut down the machine that contains the exported agent that was created during failover.
15. Perform a bare metal restore (BMR) for the source core and agent.

 **NOTE:** When you launch the restore you must use the recovery points that were imported from the target core to the agent on the virtual machine.

16. Wait for the BMR reboot and for the agent service to restart, and then view and record the network connection details of the machine.
17. Navigate to the Core Console on the source core, and, on the **Machines** tab, modify the machine protection settings to add the new network connection details.

For more information, [Configuring Machine Settings](#).

18. Navigate to the Core Console on the target core, and delete the agent from the **Replication** tab.
19. In the Core Console of the source core, set up replication again between the source and target by clicking the **Replication** tab, and then adding the target core for replication.

Reporting

About reports





Your DL Appliance lets you generate and view compliance, error, and summary information for multiple core and agent machines.

You can choose to view reports online, print reports, or export and save them in one of several supported formats. The formats from which you can choose are:

- PDF
- XLS
- XLSX
- RTF
- MHT
- HTML
- TXT
- CSV
- Image

About the reports toolbar

The toolbar available for all reports lets you print and save in two different ways. The following table describes the print and save options.

Icon	Description
	Print the report
	Print the current page
	Export a report and save it to the disk
	Export a report and show it in a new window
	Use this option to copy, paste, and e-mail the URL for others to view the report with a Web browser.

About compliance reports

Compliance Reports are available for the Core and AppAssure Agent. They provide you with a way to view the status of jobs performed by a selected core or agent. Failed jobs appear in red text. Information in the Core Compliance Report that is not associated with an agent is blank.

Details about the jobs are presented in a column view that includes the following categories:

- Core
- Protected Agent
- Type
- Summary
- Status
- Error
- Start Time
- End Time
- Time
- Total Work

About errors reports

Errors Reports are subsets of the Compliance Reports and are available for Cores and AppAssure Agents. Errors Reports include only the failed jobs listed in Compliance Reports and compile them into a single report that can be printed and exported.

Details about the errors are presented in a column view with the following categories:

- Core
- Agent
- Type
- Summary
- Error
- Start Time
- End Time
- Elapsed Time
- Total Work

About the Core Summary Report

The **Core Summary Report** includes information about the repositories on the selected Core and about the agents protected by that core. The information is displayed as two summaries within one report.

Repositories summary

The **Repositories** portion of the **Core Summary Report** includes data for the repositories located on the selected core. Details about the repositories are presented in a column view with the following categories:

- Name
- Data Path
- Metadata Path
- Allocated Space
- Used Space
- Free Space

- Compression/Dedupe Ratio

Agents summary

The **Agents** portion of the **Core Summary Report** includes data for all agents protected by the selected core.

Details about the agents are presented in a column view with the following categories:

- Name
- Protected Volumes
- Total protected space
- Current protected space
- Change rate per day (**Average, Median**)
- Jobs Statistic (**Passed, Failed, Canceled**)

Generating a report for a Core or agent

To generate a report for a core or agent:

1. Navigate to the Core Console and select the Core or the Agent for which you want to run the report.
2. Click the **Tools** tab.
3. From the **Tools** tab, expand **Reports** in the left navigation area.
4. In the left navigation area, select the report you want to run. The reports available depend on the selection you made in Step 1 and are described below.

Machine	Available Reports
Core	Compliance Report Summary Report
	Errors Report
Agent	Compliance Report Errors Report

5. In the **Start Time** drop-down calendar, select a start date, and then enter a start time for the report.



NOTE: No data is available before the time the Core or the Agent was deployed.

6. In the **End Time** drop-down calendar, select an end date, and then enter an end time for the report.
7. For a **Core Summary Report**, select the **All Time** check box if you want the **Start Time** and the **End Time** to span the lifetime of the Core.
8. For a **Core Compliance Report** or a **Core Errors Report**, use the **Target Cores** drop-down list to select the Core for which you want to view data.
9. Click **Generate Report**.


After the report generates, you can use the toolbar to print or export the report.

About the Central Management Console Core reports

Your DL Appliance lets you generate and view compliance, error, and summary information for multiple Cores. Details about the Cores are presented in column views with the same categories described in this section.

Generating a report from the Central Management Console

To generate a report from the Central Management Console:

1. From the **Central Management Console Welcome** screen, click on the drop-down menu in the upper-right corner.
2. From the drop-down menu, click **Reports** and then select one of the following options:
 - **Compliance Report**
 - **Summary Report**
 - **Failure Report**
3. From the left navigation area, select the Core or Cores for which you want to run the report.
4. In the **Start Time** drop-down calendar, select a start date, and then enter a start time for the report.
 **NOTE:** No data is available before the time the Cores are deployed.
5. In the **End Time** drop-down calendar, select an end date, and then enter an end time for the report.
6. Click **Generate Report**.

After the report generates, you can use the toolbar to print or export the report.

Getting help

Finding documentation and software updates

Direct links to AppAssure and DL1000 Appliance documentation and software updates are available from the Core Console.

Documentation

To access the link for documentation:

1. On the Core Console, click the **Appliance** tab.
2. From the left pane, navigate **Appliance** → **Documentation** link.

Software updates

To access the link for software updates:

1. On the Core Console, click the **Appliance** tab.
2. From the left pane, navigate **Appliance** → **Software Updates** link.

Contacting Dell

Dell provides several online and telephone-based support and service options. If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog. Availability varies by country and product, and some services may not be available in your area.

To contact Dell for sales, technical support, or customer-service issues, go to **software.dell.com/support**.

Documentation feedback

Click the **Feedback** link in any of the Dell documentation pages, fill up the form, and click **Submit** to send your feedback.