AppAssure 5
User Guide

Version 5.3.6

Notes, Cautions and Warnings

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

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

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

2013-11

Contents

Chapter 4: Protecting Workstations and Servers

Chapter 5: Protecting Server Clusters

1 What's New in AppAssure 5

This chapter lists new and changed features in this release of AppAssure 5.

What's New in AppAssure 5 User Guide

The following table lists the changes that are described in this version of the documentation to support AppAssure 5.

Topic	Description
Modifying AppAssure 5 Engine Settings on page 34	Modified topic. It describes how to configure the engine settings for AppAssure 5 to include parameters such as the preferable port. The recommendation to leave the No Delay option unchecked was added to the procedure; as not doing so could impact network efficiency.
Creating a Repository on page 37	Modified topic. It describes how to create a repository in AppAssure 5. The support for Windows 8.1 and Windows Server 2012 R2 was added to the sizing parameter for storage locations that are NTFS volumes.
Adding a Storage Location to an Existing Repository on page 41	Modified topic. It describes how to create a repository in AppAssure 5. The support for Windows 8.1 and Windows Server 2012 R2 was added to the sizing parameter for storage locations that are NTFS volumes.
Checking a Repository on page 44	Modified topic. It describes how to check repositories when errors occur as a result of hardware failure, improperly shutting down a server, or when a repository fails to import properly.
Configuring an Email Server and Email Notification Template on page 70	Modified topic. It describes how to configure an email server and email notification template. The steps in the procedure were updated to reflect the workflow in the user interface.

Topic	Description
About the Local Mount Utility on page 75	Modified topic. It describes the Local Mount Utility (the LMU). Updated language in this topic and in each related subtopic for clarity.
Dismounting a Recovery Point Using the Local Mount Utility on page 79	Modified topic. It describes how to dismount recovery points in the LMU. Clarified steps to launch the LMU; changed the presentation of some content to procedural steps; clarified that closing the LMU window minimizes the application; and added steps to close the LMU properly through the tray menu.
About the Local Mount Utility Tray Menu on page 80	Modified topic. It describes the LMU tray menu. Changed the presentation of content to tabular form for clarity and updated references to user interface elements.
About Protecting Workstations and Servers on page 92	Modified topic. It describes the requirements for protecting your data using the AppAssure 5 Core Console. Information about Windows 8.1 and Windows Server 2012 R2 was added to note regarding FAT32 EFI partitions. Information about the support of Microsoft Windows Storage Spaces was added. Additionally, information was added with regard to bare metal restore and Windows 8.1 Storage Spaces.
Mounting a Recovery Point for a Windows Machine on page 136	Modified topic. It describes how to mount a recovery point for a Windows machine to access stored data through a local file system. Added step regarding monitoring the completed task using the Active Task dialog box. Information about mounting recovery points from restored data with data deduplication enabled was also added to the topic.
Restoring Data on page 142	Modified topic. It describes how to export protected data from a Windows machine to a virtual machine or roll back a Windows or Linux machine to a previous recovery point. Information about Windows 8.1 and Windows Server 2012 R2 was added to note regarding FAT32 EFI partitions. Information was also added with regard to the support of Microsoft Windows Storage Spaces and data deduplication on Windows 8.1.
Exporting Backup Information for your Windows Machine to a Virtual Machine on page 144	Modified topic. It describes how to export data from a Windows machine to a virtual machine. It also lists the virtual machines that are supported. In this release of AppAssure 5, Version 5.3.6.125, a limitation exists when attempting to perform a VM export on machines that have Windows 8.1 or Windows Server 2012 R2 installed.

Topic	Description
Understanding Bare Metal Restore on page 156	Modified topic. Formerly titled About Bare Metal Restore for Windows Machines, it describes conceptual information about performing a Bare Metal Restore (BMR) for Windows and Linux machines. It includes samples of similar and dissimilar restores, and indicates BMR functionality is used for both disaster recovery and server migration. Noted that BMR not supported for Windows 8.1 Storage Spaces.
Roadmap for Performing a Bare Metal Restore for a Windows Machine on page 157	Modified topic. Expanded roadmap to include several new and updated topics. This roadmap is parallel to the new roadmap for performing BMR for Linux machines.
Prerequisites for Performing a Bare Metal Restore for a Windows Machine on page 158	Modified topic. Expanded to add image media; clarified requirement for image media and software; clarified drivers needed; added requirement for storage space and partitions. Relocated this topic.
Managing a Windows Boot Image on page 159	New topic. This topic includes the information related to the Boot CD ISO image required to perform a BMR for Windows machines.
Creating a Boot CD ISO Image for Windows on page 159	Modified topic. This topic includes steps to create a boot CD ISO image. Renamed from Creating the Boot CD.
Defining Boot CD ISO Image Parameters on page 160	New topic. This topic includes the parameters that may be required when defining a boot CD in the Create Boot CD dialog box.
Transferring the Boot CD ISO Image to Media on page 163	New topic. This topic describes burning the boot CD ISO image to media such as a CD or DVD, which is required for performing BMR on a physical machine.
Loading the Boot CD and Starting the Target Machine on page 163	Modified topic. This topic, formerly titled Loading a Boot CD, describes loading the boot CD and starting the destination machine from the boot image.
Launching a Bare Metal Restore for Windows on page 164	New topic. This topic collects the information required to launch a BMR for Windows machines.
Selecting a Recovery Point and Initiating Rollback for BMR on page 165	New topic. This topic, replacing "Launching a Restore from the AppAssure 5 Core," describes selecting the recovery point from the Core Console that contains the source of data for the BMR. It includes the step to begin the restore by connecting to the Universal Core Console. Noted that BMR not supported for Windows 8.1 Storage Spaces.

Topic	Description
Mapping Volumes for a Bare Metal Restore on page 166	Modified topic. This topic, previously entitled "Mapping Volumes," describes mapping volumes in the recovery point with those volumes on the destination machine.
Injecting Drivers to Your Target Server on page 167	New topic. This topic describes the process of injecting additional drivers to the new hardware that may be required by the restored operating system to perform properly for dissimilar hardware.
Verifying a Bare Metal Restore on page 168	New topic. This topic collects the information regarding verifying a BMR.
Starting a Restored Target Server on page 169	Modified topic. It describes how to start a server on which you have successfully completed a bare metal restore. Added a note suggesting verification of the restore, and a procedure to eject the boot CD before restarting.
Troubleshooting Connections to the Universal Recovery Console on page 170	New topic. This topic includes steps to troubleshoot connecting to the boot CD image when attempting to select a recovery point and initiate a rollback.
Repairing Startup Problems on page 170	Modified topic. This topic describes how to repair startup problems after completing a BMR. Simplified language for accuracy.
Roadmap for Performing a Bare Metal Restore on Linux Machines on page 171	New topic. This topic shows all possible steps that may be required to perform a BMR on Linux machines. This roadmap is parallel to the roadmap for performing BMR for Windows machines.
Prerequisites for Performing a Bare Metal Restore for a Linux Machine on page 172	Modified topic. This topic lists prerequisites for performing a BMR for Linux machines. Added the requirement for a suitable recovery point to restore; added similar or dissimilar hardware; changed the name of the Live CD to Live DVD; added the requirement for CD or DVD media onto which to burn the Live DVD image; and included storage drivers and network adapters. Deleted instructions to install the Screen utility, which is now included on the Live DVD.
Managing a Linux Boot Image on page 173	New topic. It includes the information related to the Live DVD boot ISO image required to perform a BMR for Linux machines.
Downloading a Boot ISO Image for Linux on page 173	New topic. It indicates the requirement for the Live DVD boot image to match the release of AppAssure 5, and describes how to download the latest Live DVD image from the license portal.

Topic	Description
Transferring the Live DVD ISO Image to Media on page 174	New topic. It describes burning the Live DVD boot ISO image to media such as a CD or DVD, which is required for performing BMR on a physical machine.
Loading the Live DVD and Starting the Target Machine on page 174	New topic. It describes loading the Live DVD boot disk and starting the destination machine from the boot image.
Managing Linux Partitions on page 175	New topic. It presents information about destination drive and what you need to consider for restoring data.
Creating Partitions on the Destination Drive on page 175	New topic. It describes how to create a partition on the destination drive of hardware used for a bare metal restore to match the volume in the recovery point.
Mounting Partitions from the Command Line on page 177	New topic. It describes creating and mounting partitions on Linux machines.
Launching a Bare Metal Restore for Linux on page 177	New topic. This topic explains two methods for performing BMR for Linux (using the Core Console UI or from the aamount utility at the command line) and lists the steps required for each.
Starting the Screen Utility on page 178	New topic. It describes the use of the optional Screen utility included on the Live DVD ISO image, which is convenient if performing a BMR from the command line.
Launching a Bare Metal Restore for a Linux Machine using the Command Line on page 179	Modified topic. Formerly entitled "Performing a Bare Metal Restore for a Linux Machine using the Command Line." This topic walks users through launching a BMR from the command line. Removed step to load the Live DVD and start the machine, and removed step to manage partitions, as these now appear in earlier sections. Added step to mount volumes before launching aamount. Moved steps for verification and creating bootable partitions to their own sections.
Verifying the Bare Metal Restore from the Command Line on page 182	New topic. It lists steps to verify a successful BMR for Linux machines.
Performing a File System Check on the Restored Volume on page 182	New topic. It indicates how to verify whether appropriate partitions are mounted, describes unmounting files, running a file system check, and remounting volumes.
Creating Bootable Partitions on the Restored Linux Machine using the Command Line on page 183	New topic. It describes in detail steps required to create bootable partitions on the restored Linux machine.

Additional Changes

This version of the AppAssure 5 User Guide also includes the following general changes:

- Editorial changes.

2 Introduction to AppAssure 5

This chapter provides an introduction and overview of AppAssure 5. It describes the features, functionality, and architecture, and consists of the following topics:

About AppAssure 5

AppAssure 5 sets a new standard for unified data protection by combining backup, replication, and recovery in a single solution that is engineered to be the fastest and most reliable backup for protecting virtual machines (VM), physical machines, and cloud environments.

AppAssure 5 combines backup and replication into one integrated and unified data protection product that also provides application awareness to ensure reliable application data recovery from your backups. AppAssure 5 is built on the new, patent-pending True Scale™ architecture which delivers the fastest backup performance with very aggressive, near-zero recovery time objectives (RTO) and recovery point objectives (RPO).

AppAssure 5 combines several unique, innovative, and breakthrough technologies:

- Live Recovery
- Recovery Assure
- Universal Recovery
- True Global Deduplication

These technologies are engineered with secure integration for cloud disaster recovery and deliver fast and reliable recovery. With its scalable object store, AppAssure 5 is uniquely capable of handling up to petabytes of data very rapidly with built-in global deduplication, compression, encryption, and replication to any private or public cloud infrastructure. Server applications and data can be recovered in minutes for data retention and compliance purposes.

Today's legacy backup tools and first generation VM backup tools are inefficient and ineffective. The outdated backup tools lack the ability to handle large-scale data and do not offer the level of performance and reliability needed for protecting business-critical applications. Combine this with complex and mixed IT environments and it presents an administrative challenge for IT professionals and vulnerability of system data.

AppAssure 5 addresses this complexity and inefficiency through our core technology and support of multi-hypervisor environments including those running on VMware vSphere and Microsoft Hyper-V, which comprise both private and public clouds. AppAssure 5 offers these technological advances while dramatically reducing IT management and storage costs.

AppAssure 5 Core Technologies

Details about the core technologies of AppAssure 5 are described in the following topics.

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. You can recover an entire volume with near-zero RTO and an RPO of minutes.

AppAssure 5 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 directly from the backup file without waiting for a full restore to production storage. Users remain productive and IT departments reduce recovery windows to meet today's increasingly stringent RTO and RPO service-level agreements.

Recovery Assure

Recovery Assure lets you perform automated recovery testing and verification of backups. It includes, but is not limited to, file systems; Microsoft Exchange 2007, 2010, and 2013; and, the different versions of Microsoft SQL Server 2005, 2008, 2008 R2, and 2012. Recovery Assure provides recoverability of applications and backups in virtual and physical environments, and features a comprehensive integrity checking algorithm based on 256-bit SHA keys that check the correctness of each disk block in the backup during archiving, replication, and data seeding operations. This ensures that data corruption is identified early and prevents corrupted data blocks from being maintained or transferred during the backup process.

Universal Recovery

Universal Recovery technology gives you unlimited machine restoration flexibility. You can restore your backups from physical to virtual, virtual to virtual, virtual to physical, or physical to physical, and carry out bare metal restores to dissimilar hardware; for example, P2V, V2V, V2P, P2P, P2C, V2C, C2P, C2V.

It also accelerates cross-platform moves among virtual machines; for instance, 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. With AppAssure 5, you can also recover or export physical to cloud, or virtual to cloud.

True Global Deduplication

AppAssure 5 provides true global deduplication that dramatically reduces your physical disk capacity requirements by offering space reduction ratios exceeding 50:1, while still meeting the data storage requirements. True Scale inline block-level compression and deduplication with line speed performance, along with built-in integrity checking, prevents data corruption from affecting the quality of the backup and archiving processes.

Product Features of AppAssure 5

Using AppAssure 5, you can manage all aspects of protection and recovery of critical data through the following features and functionality. They include:

Repository

The AppAssure repository uses deduplication volume manager (DVM) to implement a volume manager that provides support for multiple volumes, each of which could reside on different storage technologies such as Storage Area Network (SAN), Direct Attached Storage (DAS), Network Attached Storage (NAS), or cloud storage. Each volume consists of a scalable object store with deduplication. The scalable object store behaves as a records-based file system, where the unit of storage allocation is a fixed-sized data block called a record. This architecture lets you configure block-sized support for compression and deduplication. Rollup operations are reduced to metadata operations from disk intensive operations because the rollup no longer moves data but only moves the records.

The DVM can combine a set of object stores into a volume and they can be expanded by creating additional file systems. The object store files are pre-allocated and can be added on demand as storage requirements change. It is possible to create up to 255 independent repositories on a single AppAssure 5 Core and to further increase the size of a repository by adding new file extents. An extended repository may contain up to 4,096 extents that span across different storage technologies. The maximum size of a repository is 32 Exabytes. Multiple repositories can exist on a single core.

True Global Deduplication

True Global Deduplication is an effective method of reducing backup storage needs by eliminating redundant or duplicate data. Deduplication is effective because only one unique instance of the data across multiple backups is stored in the repository. The redundant data is stored, but not physically; it is simply replaced with a pointer to the one unique data instance in the repository.

Conventional backup applications have been performing repetitive full backups every week, but AppAssure performs incremental block-level backups of the machines forever. This incremental-forever approach in tandem with data deduplication helps to drastically reduce the total quantity of data committed to the disk.

The typical disk layout of a server consists of the operating system, application, and data. In most environments, the administrators often use a common flavor 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 at the same time, 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. True Global Deduplication

AppAssure 5 performs target-based inline data deduplication. This means that the snapshot data is transmitted to the Core before it is deduplicated. Inline data deduplication simply means the data is deduplicated before it is committed to disk. This is very different from at-source or post-process deduplication, where the data is deduplicated at the source before it is transmitted to the target for storage, and in post-process the data is sent raw to the target where it is analyzed and deduplicated after the data has been committed to disk. At-source deduplication consumes precious system resources on the machine whereas the post-process data deduplication approach needs all the requisite data on disk (a greater initial capacity overhead) before commencing the deduplication process. On the other hand, inline data deduplication does not require additional disk capacity and CPU cycles on the source or on the Core for the deduplication process. Lastly, conventional backup applications perform repetitive full backups every week, while AppAssure performs incremental block-level backups of the machines forever. This incremental forever approach in tandem with data deduplication helps to drastically reduce the total quantity of data committed to the disk with a reduction ratio of as much as 80:1.

Encryption

AppAssure 5 provides integrated encryption to protect backups and data-at-rest from unauthorized access and use, ensuring data privacy. AppAssure 5 provides strong encryption. By doing so, backups of protected computers are inaccessible. Only the user with the encryption key can access and decrypt the data. There is no limit to the number of encryption keys that can be created and stored on a system. DVM uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. Encryption is performed inline on snapshot data, at line speeds without impacting performance. This is because DVM implementation is multi-threaded and uses hardware acceleration specific to the processor on which it is deployed.

Encryption is multi-tenant ready. The deduplication has been specifically limited to records that have been encrypted with the same key; two identical records that have been encrypted with different keys will not be deduplicated against each other. This design decision ensures that deduplication cannot be used to leak data between different encryption domains. This is a benefit for managed service providers, as replicated backups for multiple tenants (customers) can be stored on a single core without any tenant being able to see or access other tenant data. Each active tenant encryption key creates an encryption domain within the repository where only the owner of the keys can see, access, or use the data. In a multi-tenant scenario, data is partitioned and deduplicated within the encryption domains.

In replication scenarios, AppAssure 5 uses SSL 3.0 to secure the connections between the two cores in a replication topology to prevent eavesdropping and tampering.

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.

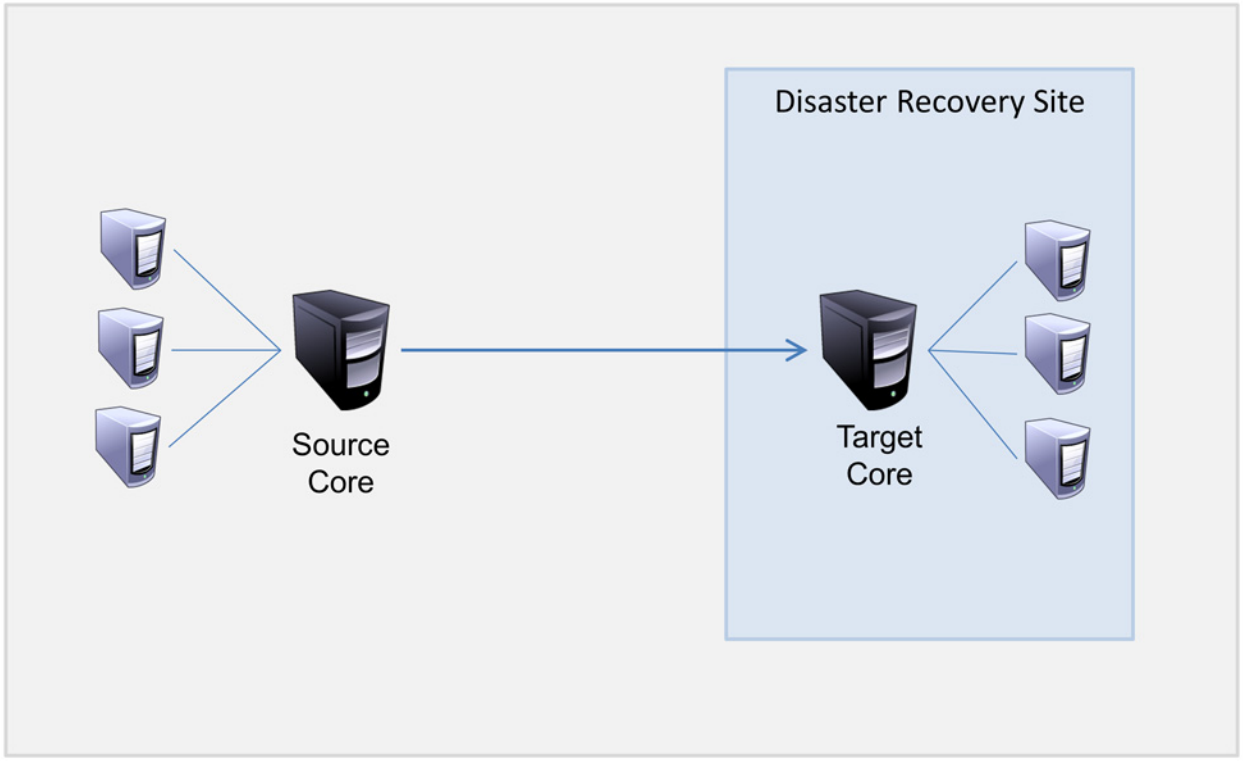

Figure 2. Replication

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 agents, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target core using external media. 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-as-a-Service (RaaS)

Managed service providers (MSPs) can fully leverage AppAssure 5 as a platform for delivering recovery as a service (RaaS). RaaS facilitates complete recovery-in-the-cloud by replicating customers' physical and virtual servers along with their data to the service provider's cloud 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.

MSPs can deploy multi-tenant AppAssure 5-based RaaS infrastructure that can host multiple and discrete organizations or business units (the tenants) that ordinarily do not share security or data on a single server or a group of servers. The data of each tenant is isolated and secure from other tenants and the service provider.

Retention and Archiving

AppAssure 5 offers flexible backup and retention policies that are easily configurable. The ability to tailor retention polices to the needs of an organization not only helps to meet compliance requirements but does so without compromising recovery time objectives (RTO).

Retention policies enforce the periods of time in 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 non-compliance data. The archive feature supports extended retentions for compliance and non-compliance data, as well as being used for seeding replication data to a target core.

Figure 3. Retention Policy

In AppAssure 5 retention policies can be customized to specify the length of time a backup recovery point is maintained. As the age of the recovery points approach the end of their retention period, they age out and are removed from the retention pool. Typically, this process becomes inefficient and eventually fails as the amount of data and the period of retention start growing rapidly. AppAssure 5 solves the big data problem by managing the retention of large amounts of data with complex retention policies and performing rollup operations for aging data using efficient metadata operations.

Backups can be performed with an interval of a few minutes; and, these backups age over days, months, and years. Retention policies manage the aging and deletion of old backups. A simple waterfall method defines the aging process. The levels within the waterfall are defined in minutes, hours, and days; weeks, months, and years. The retention policy is enforced by the nightly rollup process.

For long term archival, AppAssure 5 lets you create an archive of the source or target core on any removable media. The archive is internally optimized, and all data in the 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 will span across multiple devices based on the available space on the media. Recovery from an archive does not require a new core; any core can ingest the archive and recover data if the administrator has the passphrase and the encryption keys.

Virtualization and Cloud

The AppAssure 5 Core is cloud-ready, which allows you to leverage the compute capacity of the cloud for recovery.

AppAssure 5 can export any protected or replicated machine to licensed versions of VMware or Hyper-V. Exports can be ad-hoc or continuous. With continuous exports, the virtual machine is incrementally updated after every snapshot. The incremental updates are very fast and provide standby-clones that are ready to be powered up with a click of a button. The supported exports are VMware Workstation/Server on a folder, direct export to a Vsphere / VMware ESX(i) host, and Microsoft Server 2008 R2 Hyper-V.

Alerts and Event Management

In addition to HTTP REST APIs, AppAssure 5 also includes an extensive set of features for event logging and notification using email, syslog, or Windows Event Log. Email notifications can be used to alert users or groups of the health or status of different events in response to an alert. The syslog and Windows Event Log methods are used for centralized logging to a repository in multi-operating system environments; while in Windows-only environments, only the Windows Event Log is used.

AppAssure 5 License Portal

The AppAssure 5 License Portal provides easy-to-use tools for managing license entitlements. You can download, activate, view, and manage license keys and create a company profile to track your license assets. Additionally, the portal enables service providers and re-sellers to track and manage their customer licenses.

Web Console

AppAssure 5 features a Web-based central console that manages distributed AppAssure 5 cores from one central location. MSPs and enterprise customers with multiple distributed cores can deploy this console to get a unified view for centralized management. The AppAssure 5 Central Management Console lets you organize the managed cores in hierarchical organizational units. These organizational units can represent business units, locations, or customers for MSPs with role-based access. Using the central console, you can also run reports across all of your managed cores.

Service Management APIs

AppAssure 5 comes bundled with a set of service management APIs and provides programmatic access to all of the functionality available through the AppAssure 5 Central Management Console. The service management API is a REST API. All the API operations are performed over SSL and are mutually authenticated using X.509 v3 certificates. The management service can be accessed from within the environment or directly over the Internet from any application that can send and receive an HTTPS request and response. The approach facilitates easy integration with any Web application such as relationship management methodology (RMM) tools or billing systems. Also included with AppAssure 5 is an SDK client for PowerShell scripting.

White Labeling

AppAssure 5 can be re-branded and white-labeled for select enterprise and OEM partners under the Platinum service provider program. The Platinum service provider program lets partners brand AppAssure 5 with their custom name, logo, and color themes and deliver the product or service with their own branding and look-and-feel to their customers.

As an AppAssure partner, you can tailor the software to meet your business requirements. To further explore how you can brand AppAssure 5 to suit your business needs, contact AppAssure Sales at sales@appassure.com for more information.

This page is intentionally left blank.

3 Working with the AppAssure 5 Core

This chapter describes the various aspects of working with, configuring, and managing the AppAssure 5 Core. It includes the following topics:

Accessing the AppAssure 5 Core Console

Complete the following steps to access the AppAssure 5 Core Console.

To access the AppAssure 5 Core Console

1. Perform one of the following to access the AppAssure 5 Core Console:
   a. Log on locally to your AppAssure 5 Core server, and then select the **Core Console** icon.
   b. Or, type one of the following URLs in your Web browser:
      - https://<yourCoreServerName>:8006/apprecovery/admin/core, or
      - https://<yourCoreServerIPaddress>:8006/apprecovery/admin/core

Roadmap for Configuring the AppAssure 5 Core

Before you can use AppAssure 5, you must configure the AppAssure 5 Core. Configuration includes tasks such as creating and configuring the repository for storing backup snapshots, defining encryption keys for securing protected data, and setting up alerts and notifications. After you complete the configuration of the AppAssure 5 Core, you can then protect agents and perform recovery.

Configuring the AppAssure 5 Core involves understanding certain concepts and performing the following initial operations:

- **Create a repository**. For more information about repositories, see "Creating a Repository" on page 37.
- **Configure encryption keys**. For more information on configuring encryption keys, see "Adding an Encryption Key" on page 45.
- **Configure event notification**. For more information on configuring event notifications, see "Configuring an Email Server and Email Notification Template" on page 70.
- **Configure retention policy**. For more information on configuring retention policies, see "Managing Retention Policies" on page 81.
- **Configure SQL attachability**. For more information on configuring SQL attachability, see "Configuring SQL Attachability Settings" on page 84.

Managing Licenses

AppAssure 5 lets you can manage AppAssure 5 licenses directly from the AppAssure 5 Core Console. From the console, you can change the license key and contact the license server. You can also access the AppAssure 5 License Portal from the Licensing page in the console.

The Licensing page includes the following information:

- License type
- License status
- License pool size
- 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

For more information, see the documentation, Managing AppAssure 5 Licenses, on the AppAssure 5 Technical Documentation page at: http://docs.appassure.com/display/AA50D/AppAssure+5+Technical+Documentation.

Changing a License Key

Complete the steps in this procedure to change a license key from within the AppAssure 5 Core Console.

For information about obtaining a license key, see the documentation, Managing AppAssure 5 Licenses, on the AppAssure 5 Technical Documentation page at: http://docs.appassure.com/display/AA50D/AppAssure+5+Technical+Documentation

To change a license key

1. Navigate to the AppAssure 5 Core Console and then select the Configuration tab.
2. Click **Licensing**.

   The Licensing page appears.
3. From the license details, click **Change.**

   The Change License Key dialog box appears.
4. In the Change License Key dialog box, enter the new license key and click **OK.**

Contacting the License Portal Server

The AppAssure 5 Core Console frequently contacts the portal server to remain current with any changes made in the license portal. Typically, communication with the portal server occurs automatically at designated intervals; however, you can initiate communication on-demand.

Complete the steps in this procedure to contact the portal server.

To contact the portal server

1. Navigate to the AppAssure 5 Core Console and then click the Configuration tab.

2. Click **Licensing.**

   The Licensing page appears.

3. From the License Server option, click **Contact Now.**

Managing AppAssure 5 Core Settings

The AppAssure 5 Core settings are used to define various settings for configuration and performance. 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

Complete the steps in this procedure to change the Core display name.

To change the Core display name

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the General area, click **Change**.

   The Display Name dialog box displays.

3. In the Name text box, enter a new display name for the Core.
4. Click **OK**.

Adjusting the Nightly Job Time

Complete the steps in this procedure to adjust the nightly job time.

To adjust the nightly job time

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Nightly Jobs area, click **Change**.

   The Nightly Jobs dialog box displays.

3. In the Start 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 the maximum number of retries for transferring data.

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

To modify the transfer queue settings

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Transfer Queue area, click **Change**.

   The Transfer Queue dialog box displays.
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 on network and other system resources. As the number of agents that are 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 Timeout Settings

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

To adjust the client timeout settings

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Client Timeout Settings Configuration area, click **Change**.

   The Client Timeout Settings dialog box displays.
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 you want to lapse before a time out occurs during a read/write event.
5. Click **OK**.

Configuring Deduplication Cache Settings

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

To configure deduplication cache settings

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Deduplication Cache Configuration area, click **Change**.

   The Deduplication Cache Configuration dialog box displays.
3. In the Primary Cache Location text box, enter an updated value to change the primary cache location.
4. In the Secondary Cache Location text box, enter an updated value to change the secondary cache location.
5. In the Metadata Cache Location text box, enter an updated value to change the metadata cache location.
6. Click **OK**.

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

Modifying AppAssure 5 Engine Settings

Complete the steps in this procedure to modify AppAssure 5 engine settings.

To modify AppAssure 5 engine settings

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Replay Engine Configuration area, click **Change**.

   The Replay Engine Configuration dialog box displays.
3. Enter the configuration information as described in the following table.

Text Box	Description
IP Address	Specify the IP address by choosing one of the following: □ Click **Automatically Determined** to use the preferred IP address from your TCP/IP. □ Or, click **Use a specific address** to manually enter an IP address.
Preferable Port	Enter a port number or accept the default setting. The default port is 8007. The port is used to specify the communication channel for the AppAssure engine.
Port in use	Represents the port that is in use for the Replay Engine configuration.

Text Box	Description
Allow port auto-assigning	Click for allow for automatic TCP port assignment.
Admin Group	Enter a new name for the administration group. The default name is BUILTIN\Administrators.
Minimum 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.
No Delay	It is recommended that you leave this check box unchecked as doing otherwise will impact network efficiency. If you determine that you need to modify this setting, contact Dell Support for guidance in doing so.

4. Click **OK**.

Modifying Database Connection Settings

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

To modify database connection settings

1. Navigate to the AppAssure 5 Core Console and click the Configuration tab, and then **Settings**.
2. In the Database Connection Settings area, perform one of the following:
   - Click **Apply Default**.
   - Or, click **Change**.

   The Database Connection Settings dialog box displays.
3. Enter the settings for modifying the database connection as described in the following table.

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 logon credentials for accessing the database connection.

Text Box	Description
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**.

About Repositories

A repository is used to store the snapshots that are captured from your protected workstations and servers. The repository can reside on different storage technologies such as Storage Area Network (SAN), Direct Attached Storage (DAS), or Network Attached Storage (NAS).

When you create a repository, the AppAssure 5 Core pre-allocates the storage space required for the data and metadata in the specified location. You can create up to 255 independent repositories on a single core that span across different storage technologies; in addition, you can further increase the size of a repository by adding new file extents or specifications. An extended repository can contain up to 4096 extents that span across different storage technologies.

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.

  NOTE: AppAssure 5 repositories should be stored on primary storage devices. Archival storage devices such as Data Domain are not supported due to performance limitations. Similarly, repositories should not be stored on NAS filers that tier to the cloud as these devices tend to have performance limitations when used as primary storage.

Roadmap for Managing a Repository

Before you can use AppAssure 5, you need to set up one or more repositories on the AppAssure 5 Core server. A repository stores your protected data; more specifically, it stores the snapshots that are captured from the protected servers in your environment.

When you configure a repository, you can perform a variety of tasks such as specifying where to locate the data storage on the Core server, how many locations should be added to each repository, the name of the repository, how many current operations the repositories support, and so on.

When you create a repository, the Core pre-allocates the space required for storing data and metadata in the specified location. You can create up to 255 independent repositories on a single core. To further increase the size of a single repository, you can add new storage locations or volumes.

Managing a repository involves creating, configuring, and viewing a repository and includes the following operations:

- **Access the Core Console.** For more information on how to access the AppAssure 5 Core Console, see "Accessing the AppAssure 5 Core Console" on page 30.
- **Create a repository.** For more information about creating a repository, see "Creating a Repository" on page 37.
- **View repository details.** For more information about viewing repository details, see "Viewing Details about a Repository" on page 40.
- **Modify repository settings.** For more information about modifying repository settings, see "Modifying Repository Settings" on page 40.
- **Add a new storage location.** For more information on adding a new storage location, see "Adding a Storage Location to an Existing Repository" on page 41.
- **Check a repository.** For more information about checking a repository, see "Checking a Repository" on page 44.
- **Delete a repository.** For more information about deleting a repository, see "Deleting a Repository" on page 44.

Creating a Repository

Complete the following steps to create a repository.

To create a repository

1. From the AppAssure 5 Core Console, click the Configuration tab.

   The Repositories page displays.

2. In the **Actions** drop-down menu, click **Add New Repository**.

   The Add New Repository dialog box displays.

3. Enter the information as described in the following table.

Text Box	Description
Repository Name	Enter the display name of the repository. By default, this text box consists of the word Repository and an index number, which corresponds to the number of the new repository. You can change the name as needed. You can enter up to 150 characters.
Concurrent Operations	Define the number of concurrent requests you want the repository to support. By default the value is 64.
Comments	Optionally, enter a descriptive note about this repository.

4. Click **Add Storage Location** to define the specific storage location or volume for the repository.

If the AppAssure repository that you are creating in this step is later removed, all files at the storage location of your repository will be deleted. If you do not define a dedicated folder to store the repository files, then those files will be stored in the root; deleting the repository will also delete the entire contents of the root, resulting in catastrophic data loss.

NOTE: AppAssure 5 repositories should be stored on primary storage devices. Archival storage devices such as Data Domain are not supported due to performance limitations. Similarly, repositories should not be stored on NAS filers that tier to the cloud as these devices tend to have performance limitations when used as primary storage.

The Add Storage Location dialog box appears.

5. Specify how to add the file for the storage location. You can choose to add the file on the local disk or on CIFS share.
   - Select **Add file on local disk** to specify a local machine, and then enter the information as described in the following table.

Text Box	Description
Metadata Path	Enter the location for storing the protected metadata. For example, type `X:\Repository\Metadata`. 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.
Data Path	Enter the location for storing the protected data. For example, type `X:\Repository\Data`. The same limitations to the path apply; use only alphanumeric characters, hyphen, or period, with no spaces or special characters.

   - Or, select **Add file on CIFS share** to specify a network share location, and then enter the information as described in the following table.

Text Box	Description
UNC Path	Enter the path for the network share location. If this location is at the root, define a dedicated folder name (for example, `Repository`). The path must begin with \\. When specifying the path, use only alphanumeric characters, the hyphen, and the period (only to separate host names and domains). The letters a to z are case-insensitive. Do not use spaces. No other symbols or punctuation characters are permitted.
User Name	Specify a user name for accessing the network share location.
Password	Specify a password for accessing the network share location.

6. In the Details pane, click **Show/Hide Details** and enter the details for the storage location as described in the following table.

Text Box	Description
Size	Set the size or capacity for the storage location. The default is 250 MB. You can choose from the following: □ MB □ GB □ TB **NOTE:** The size that you specify cannot exceed the size of the volume. If the storage location is a New Technology File System (NTFS) volume using Windows XP or Windows 7, the file size limit is 16 TB. If the storage location is a NTFS volume using Windows 8, 8.1 or Windows Server 2012, 2012 R2, the file size limit is 256 TB. **NOTE:** For AppAssure 5 to validate the operating system, Windows Management Instrumentation (WMI) must be installed on the intended storage location.
Write Caching Policy	The write caching policy controls how the Windows Cache Manager is used in the repository and helps to tune the repository for optimal performance on different configurations. Set the value to one of the following: □ On □ Off □ Sync If set to On, which is the default, Windows controls the caching. **NOTE:** Setting the write caching policy to On could result in faster performance. If you are using a version of Windows Server prior to Server 2012, the recommended setting is Off. If set to Off, AppAssure 5 controls the caching. If set to Sync, Windows controls the caching as well as the synchronous input/output.
Bytes per Sector	Specify the number of bytes you want each sector to include. The default value is 512.
Average Bytes per Record	Specify the average number of bytes per record. The default value is 8192.

7. Click **Save**.

   The Repositories screen displays to include the newly added storage location.

8. Optionally, repeat Step 4 through Step 7 to add additional storage locations for the repository.

9. Click **Create** to create the repository.

   The Repository displays in the Configuration tab.

Viewing Details about a Repository

Complete the following step to view the details for a repository.

To view details about a repository

- In the AppAssure 5 Core Console, click the Configuration tab.

  The Repositories page displays.

- Click the right angle bracket **>** symbol next to the Status column of the repository for which you want to view the details.

- From the expanded view, you can perform the following actions for a repository:
  - Modify Settings
  - Add a Storage Location
  - Check a Repository
  - Delete a Repository

- Details also display for the repository to include the storage locations and statistics. Storage location details include:
  - Metadata Path
  - Data Path
  - Size

  The statistical information available for you to view consist of:
  - Deduplication
  - Record I/O
  - Storage Engine

  The level of detail available for Deduplication is reported as the number of block dedupe hits, block dedupe misses, and block compression rate.

  The detail rendered for Record I/O consists of the rate (MB/s), read rate (MB/s), and write rate (MB/s).

  The storage engine details are include the rate (MB/s), read rate (MB/s), and write rate (MB/s),

Modifying Repository Settings

After you add a repository, you can modify the repository settings such as the description or the maximum concurrent operations. You can also add a new storage location for the repository. For more information on adding a new storage location, see "Adding a Storage Location to an Existing Repository" on page 41.

To modify repository settings

1. In the AppAssure 5 Core Console, click the Configuration tab.

   The Repositories page displays.

2. Click the right angle bracket **>** symbol next to the Status column of the repository that you want to modify.
3. Next to **Actions**, click **Settings**.

   The Repository Settings dialog box displays.
4. Edit the repository information as described in the following table.

Text Box	Description
Repository Name	Represents the display name of the repository. By default, this text box consists of the word Repository and an index number, which corresponds to the number of the repository. **NOTE:** You cannot edit the repository name.
Description	Optionally, enter a descriptive note about the repository.
Maximum Concurrent Operations	Define the number of concurrent requests you want the repository to support.
Enable Deduplication	Clear this checkbox to turn off deduplication, or select this checkbox to enable deduplication. **NOTE:** Changing this setting only applies to backups taken after the setting has been made. Existing data, or data replicated from another core or imported from an archive, will retain the deduplication values in place at the time the data was captured from an agent.
Enable Compression	Clear this checkbox to turn off compression, or select this checkbox to enable compression. **NOTE:** This setting applies only to backups taken after the setting has been changed. Existing data, or data replicated from another core or imported from an archive, will retain the compression values in place at the time the data was captured from an agent.

5. Click **Save**.

Adding a Storage Location to an Existing Repository

Adding a storage location lets you define where you want the repository or volume to be stored. Complete the steps in the following procedure to specify the storage location for the repository or volume.

To add a storage location to an existing repository

1. Click the right angle bracket **>** symbol next to the Status column of the repository for which you want to add a new storage location.
2. In the **Actions** pane, click **Add Storage Location.**

   The Add Storage Location dialog box displays.
3. Specify how to add the file for the storage location. You can choose to add the file on the local disk or on CIFS share.

- Select **Add file on local disk** to specify a local machine and then enter the information as described in the following table.

Text Box	Description
Metadata Path	Enter the location for storing the protected metadata. For example, type `X:\Repository\Metadata`. 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.
Data Path	Enter the location for storing the protected data. For example, type `X:\Repository\Data`. The same limitations to the path apply; use only alphanumeric characters, hyphen, or period, with no spaces or special characters.

- Or, select **Add file on CIFS share** to specify a network share location and then enter the information as described in the following table.

Text Box	Description
UNC Path	Enter the path for the network share location. If this location is at the root, define a dedicated folder name (for example, `Repository`). The path must begin with \\. When specifying the path, use only alphanumeric characters, the hyphen, and the period (only to separate host names and domains). The letters a to z are case-insensitive. Do not use spaces. No other symbols or punctuation characters are permitted.
User Name	Specify a user name for accessing the network share location.
Password	Specify a password for accessing the network share location.

4. In the Details pane, click **Show/Hide Details** and enter the details for the storage location as described in the following table.

Text Box	Description
Size	Set the size or capacity for the storage location.The default size is 250 MB. You can choose from the following: □ MB □ GB □ TB **NOTE:** The size that you specify cannot exceed the size of the volume. If the storage location is a New Technology File System (NTFS) volume using Windows XP or Windows 7, the file size limit is 16 TB. If the storage location is a NTFS volume using Windows 8, 8.1 or Windows Server 2012, 2012 R2, the file size limit is 256 TB. **NOTE:** For AppAssure 5 to validate the operating system, Windows Management Instrumentation (WMI) must be installed on the intended storage location.
Write Caching Policy	The write caching policy controls how the Windows Cache Manager is used in the repository and helps to tune the repository for optimal performance on different configurations. Set the value to one of the following: □ On □ Off □ Sync If set to On, which is the default, Windows controls the caching. **NOTE:** Setting the write caching policy to On could result in faster performance. If you are using a version of Windows Server prior to Server 2012, the recommended setting is Off. If set to Off, AppAssure 5 controls the caching. If set to Sync, Windows controls the caching as well as the synchronous input/output.
Bytes per Sector	Specify the number of bytes you want each sector to include. The default value is 512.
Average Bytes per Record	Specify the average number of bytes per record. The default value is 8192.

5. Click **Save**.

   The Repositories screen displays to include the newly added storage location.

6. Optionally, repeat Step 4 through Step 7 to add additional storage locations for the repository.

7. Click **OK**.

Checking a Repository

AppAssure 5 provides the ability to perform a diagnostic check of a repository volume when errors occur. Core errors could be the result of it being improperly shut down, a hardware failure, and so on.

This procedure should only be performed for diagnostic purposes; for example, in the event of hardware failure, improper shutdown of the Core, failure when importing a repository, and so on.

To check a repository

1. On the Configuration tab, click **Repositories**, and then select the right angle bracket **>** symbol next to the repository you want to check.
2. In the **Actions** pane, click **Check**.

   The Check Repository dialog box appears.
3. In the Check Repository dialog box, click **Check**.

   NOTE: If the check fails, you will need to restore the repository from an archive.

Deleting a Repository

Complete the steps in this procedure to delete a repository.

To delete a repository

1. On the Configuration tab, click **Repositories**, and then select the right angle bracket **>** symbol next to the repository you want to delete.
2. In the **Actions** pane, click **Delete**.
3. In the Delete Repository dialog box, click **Delete**.

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

Managing Security

The AppAssure 5 Core can encrypt agent snapshot data within the repository. Instead of encrypting the entire repository, AppAssure 5 lets you specify an encryption key during the protection of an agent in a repository, which allows the keys to be re-used for different agents. Encryption does not impact performance, as each active encryption key creates an encryption domain. Thus letting a single core support multi-tenancy by hosting multiple encryption domains. In a multi-tenant environment, data is partitioned and deduplicated within the encryption domains. Because you manage the encryption keys, loss of the volume cannot leak the keys.

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 encryption keys that are configured on the AppAssure 5 Core.
- There is no limit to the number of encryption keys you can create on the Core.

Adding an Encryption Key

Complete the steps in this procedure to add an encryption key.

To add an encryption key

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option on the Configuration tab, select **Security**.
3. Click **Actions**, and then click **Add Encryption Key**.

   The Create Encryption Key dialog box appears.
4. In the Create Encryption Key dialog box, enter the details for the key as described in the following table.

Text Box	Description
Name	Enter a name for the encryption key.
Description	Enter a comment for the encryption key. It is used to provide additional 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.

5. Click **OK**.

AppAssure 5 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.

Editing an Encryption Key

Complete the steps in this procedure to edit an encryption key.

To edit an encryption key

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Security**.

   The Encryption Keys screen displays.
3. Click right angle bracket **>** symbol next to the name of the encryption key that you want to edit, and then click **Edit**.

   The Edit Encryption Key dialog box displays.
4. In the Edit Encryption Key dialog box, edit the name or modify the description for the encryption key.
5. Click **OK**.

Changing an Encryption Key Passphrase

Complete the steps in this procedure to change the passphrase for an encryption key.

To change an encryption key passphrase

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Security**.
3. Click right angle bracket **>** symbol next to the name of the encryption key you want to edit, and then click **Change Passphrase**.

   The Change Passphrase dialog box displays.
4. In the Change Passphrase dialog box, enter the new passphrase for the encryption and then re-enter the passphrase to confirm what you entered.
5. Click **OK**.

AppAssure 5 uses AES 256-bit encryption in the Cipher Block Chaining (CBC) mode with 256-bit keys. It is recommended 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.

Importing an Encryption Key

Complete the steps in this procedure to import an encryption key.

To import an encryption key

1. Navigate to the AppAssure 5 Core Console, and then click the Configuration tab.
2. From the Manage option, click **Security**.

3. Click the **Actions** drop-down menu, and then click **Import**.

   The Import Key dialog box displays.

4. In the Import Key dialog box, click **Browse** to locate the encryption key you want to import, and then click **Open**.

5. Click **OK**.

Exporting an Encryption Key

Complete the steps in this procedure to export an encryption key.

To export an encryption key

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.

2. From the Manage option, click **Security**.

3. Click right angle bracket **>** symbol next to the name of the encryption key you want to export, and then click **Export**.

   The Export Key dialog box displays.

4. In the Export Key dialog box, click **Download Key** to save and store the encryption keys in a secure location.

5. Click **OK**.

Removing an Encryption Key

Complete the steps in this procedure to remove an encryption key.

To remove an encryption key

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.

2. From the Manage option, click **Security**.

3. Click right angle bracket **>** symbol next to the name of the encryption key you want to remove, and then click **Remove**.

   The Remove Key dialog box displays.

4. In the Remove Key dialog box, click **OK** to remove the encryption key.

   NOTE: Removing an encryption key does not make the data un-encrypted.

Understanding Replication

This section provides conceptual and procedural information to help you understand and configure replication in AppAssure 5.

About 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. The source core copies the recovery points of the protected agents and then asynchronously and continuously transmits them to a target core at a remote disaster recovery site. The off-site location can be a company-owned data center (self-managed core) or a third-party managed service provider's (MSP's) location or cloud environment. When replicating to a MSP, you can use built-in work flows that let you request connections and receive automatic feedback notifications.

Possible scenarios for replication include:

- **Replication to a Local Location.** The target core is located in a local data center or on-site location, and replication is maintained at all times. In this configuration, the loss of the Core would not prevent a recovery.
- **Replication to an Off-site Location.** The target core is located at an off-site disaster recovery facility for recovery in the event of a loss.
- **Mutual Replication.** Two data centers in two different locations each contain a core and are protecting agents and serving as the off-site disaster recovery backup for each other. In this scenario, each core replicates the agents to the Core that is located in the other data center.
- **Hosted and Cloud Replication.** AppAssure MSP partners maintain multiple target cores in a data center or a public cloud. On each of these cores, the MSP partner lets one or more of their customers replicate recovery points from a source core on the customer's site to the MSP's target core for a fee.

  NOTE: In this scenario, customers would only have access to their own data.

Possible replication configurations include:

- **Point to Point**. Replicates a single agent from a single source core to a single target core.

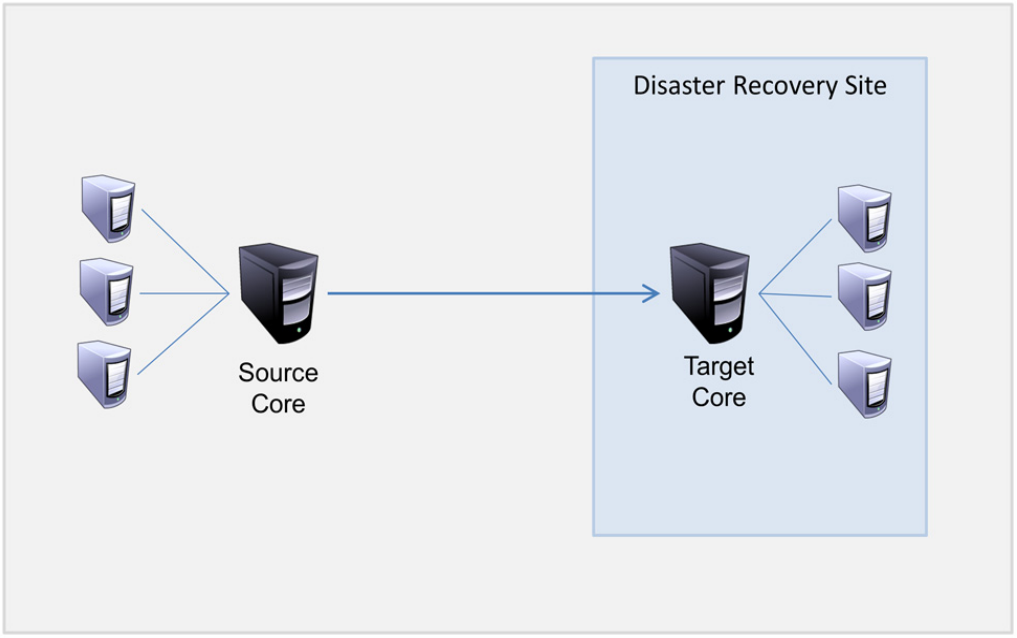

Figure 4. Point to point configuration

- **Multi-Point to Point.** Replicates multiple source cores to a single target core.

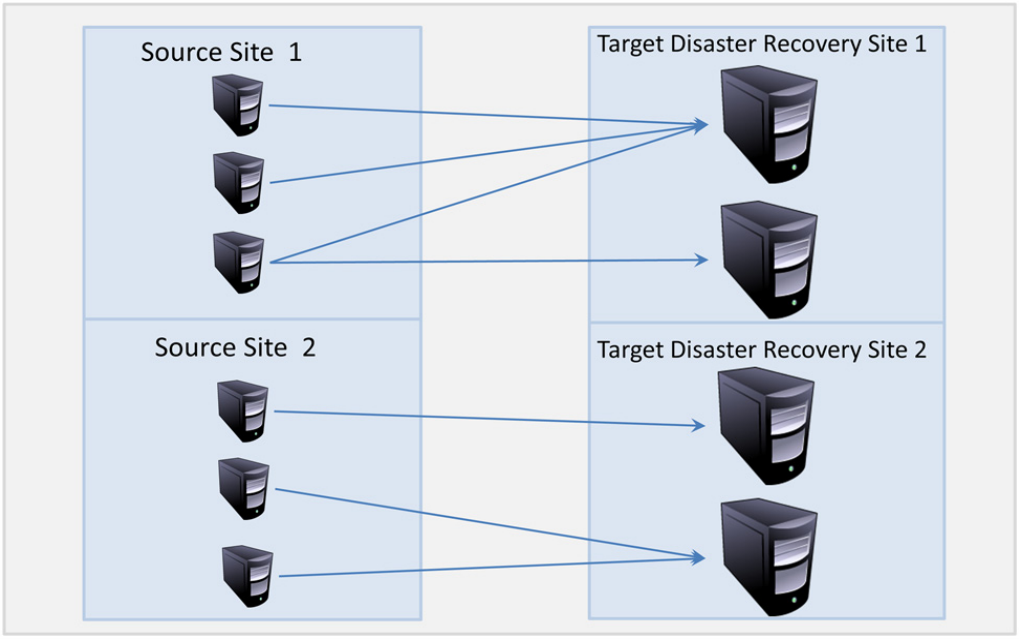

Figure 5. Multi-point to point configuration

About Seeding

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 using external media to transfer the initial data to the target core. This is typically useful for large sets of data or sites with slow links.

While it is possible to seed the base data over a network connection, it is not recommended. Initial seeding involves potentially very large amounts of data, which could overwhelm a typical WAN connection. For example, if the seed data measures 10 GB and the WAN link transfers 24 Mbps, the transfer could take more than 40 days to complete.

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 are replicated to the target site. After the target core consumes the seeding archive, the newly replicated incremental recovery points automatically synchronize.

Seeding is a two-part process (also known as copy-consume):

- The first part involves copying, which is the writing of the initial replicated data to a removable media source. Copying duplicates all of the existing recovery points from the source core to a local removable storage device such as a USB drive. After copying is complete, you must then transport the drive from the source core location to the remote target core location.
- The second part is consuming, which occurs when a target core receives the transported drive and copies the replicated data to the repository. The target core then consumes the recovery points and uses them to form replicated agents.

While replication of incremental snapshots can occur between the source and target cores before seeding is complete, the replicated snapshots transmitted from the source to the target will remain "orphaned" until the initial data is consumed, and they are combined with the replicated base images. For more information about orphaned recovery points, see "Deleting an Orphaned Recovery Point Chain" on page 140.

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.

About Failover and Failback in AppAssure 5

In the case of a severe outage in which your source core and agents fail, AppAssure 5 supports failover and failback in replicated environments. 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 that was protected by the failed source core. 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. For more information, see "Roadmap for Failover and Failback" on page 65.

About Replication and Encrypted Recovery Points

While the seed drive does not contain backups of the source core registry and certificates, the seed drive does contain encryption keys from the source core if the recovery points being replicated from source to target are encrypted. The replicated recovery points remain encrypted after they are transmitted to the target core. The owners or administrators of the target core need the passphrase to recover the encrypted data.

About Retention Policies for Replication

The retention policy on the source core determines the retention policy for the data replicated to the target core, because the replication task transmits the merged recovery points that result from a rollup or ad-hoc deletion.

For more information on retention policies, see "Managing Retention Policies" on page 81.

The target core is not capable of rollup or of ad-hoc deletion of recovery points. These actions can only be performed by the source core.

Performance Considerations for Replicated Data Transfer

If the bandwidth between the source core and the target core cannot accommodate the transfer of stored recovery points, replication begins with seeding the target core with base images and recovery points from the selected servers protected on the source core. The seeding process only has to be performed once, as it serves as the foundation that is required for regularly scheduled replication.

When preparing for replication, you should consider the following factors:

- **Change Rate**. The change rate is the rate at which the amount of protected data is accumulated. The rate depends on the amount of data that changes on protected volumes and the protection interval of the volumes. If a set of blocks change on the volume, reducing the protection interval reduces the change rate.
- **Bandwidth**. The bandwidth is the available transfer speed between the source core and the target core. It is crucial that the bandwidth be greater than the change rate for replication to keep up with the recovery points created by the snapshots. Due to the amount of data transmitted from core to core, multiple parallel streams may be required to perform at wire speeds up to the speed of a 1GB Ethernet connection.

  NOTE: Bandwidth specified by the ISP is the total available bandwidth. The outgoing bandwidth is shared by all devices on the network. Make sure that there is enough free bandwidth for replication to accommodate the change rate.
- **Number of Agents**. It is important to consider the number of agents protected per source core and how many you plan to replicate to the target. AppAssure 5 lets you perform replication on a per-protected server basis, so you can choose to replicate certain servers. If all protected servers must be replicated, this drastically affects the change rate, particularly if the bandwidth between the source and target cores is insufficient for the amount and size of the recovery points being replicated.

Depending on your network configuration, replication can be a time-consuming process.

The Maximum Change Rate for WAN Connection Types is shown in the table below with examples of the necessary bandwidth per gigabyte for a reasonable change rate.

Broadband	Bandwidth	Max Change Rate
DSL	768 Kbps and up	330 MB per hour
Cable	1 Mbps and up	429 MB per hour
T1	1.5 Mbps and up	644 MB per hour
Fiber	20 Mbps and up	8.38 GB per hour

For optimum results, you should adhere to the recommendations listed in the table above.

If a link fails during data transfer, replication resumes from the previous failure point of the transfer once link functionality is restored.

Roadmap for Performing Replication

To replicate data using AppAssure 5, you must configure the source and target cores for replication. After you configure replication, you can then replicate agent data, monitor and manage replication, and perform recovery.

Performing replication in AppAssure 5 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" on page 53.
- **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" on page 55.
- **Replicate a new agent attached to the source core.** For more information on replicating an agent, see "Replicating a New Agent" on page 116.
- **Replicate an existing agent.** For more information on configuring an agent for replication, see "Replicating Agent Data on a Machine" on page 118.
- **Consume the Seed Drive.** For more information on consuming seed drive data on the target core, see "Consuming the Seed Drive on a Target Core" on page 59.
- **Set replication priority for an agent.** For more information on prioritizing the replication of agents, see "Setting Replication Priority for an Agent" on page 119.
- **Monitor replication as needed.** For more information on monitoring replication, see "Monitoring Replication" on page 60.
- **Manage replication settings as needed.** For more information on managing replication settings, see "Managing Replication Settings" on page 62.
- **Recover replicated data in the event of disaster or data loss.** For more information on recovering replicated data, see "Recovering Replicated Data" on page 64.

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.

This configuration applies to Replication to an Off-site Location and to Mutual Replication. The AppAssure 5 Core must be installed on all source and target machines. If you are configuring AppAssure 5 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

Complete the steps in the following procedure to configure your source core to replicate to a self-managed target core.

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

1. Navigate to the AppAssure 5 Core, and then click the Replication tab.
2. In the **Actions** drop-down menu, click **Add Remote Core**.

   The Select Replication Type dialog box appears.
3. Select **I have my own remote core I wish to replicate to**, 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.
Port	Enter the port number on which the AppAssure 5 Core will communicate 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.

4. Click **Continue**.
5. In the Add Remote Core dialog box, select one of the following options:

Option	Description
Replace an existing replicated Core	Replaces an existing core on the remote host with the Core selected from the drop-down list.
Create a new replicated Core on <host name>	Creates a new core with the name in the text box on the remote target core machine. **NOTE:** This is the default selection. The core name automatically appears in the text box.

6. Select the agents you want to replicate, and then select a repository for each agent.
7. If you plan to perform the seeding process for the transfer of the base data, select the check box next to **Use a seed drive to perform initial transfer.**
8. Click **Start Replication**.
   - If you selected the option, Use a seed drive to perform initial transfer, the Copy to Seed Drive dialog box appears.
   - If you did not select to use a seed drive, the task is complete.

9. In the Copy to Seed Drive dialog box, enter the information described in the following table.

Text Box	Description
Location	Enter the path to the drive on which you want to save the initial data, such as a local USB drive.
User Name	Enter the user name for connecting to the drive. **NOTE:** This is required if the seed drive is located on a network share.
Password	Enter the password for connecting to the drive. **NOTE:** This is required if the seed drive is located on a network share.
Maximum Size	Select one of the following options: □ The entire target □ A portion of the drive's available space. Then, to designate a portion of the drive, do the following: **a.** Enter the amount of space in the text box. **b.** Select the measurement.
Recycle action	In the event the path already contains a seed drive, select one of the following options: □ **Do not reuse**. Does not overwrite or clear any existing data from the location. If the location is not empty, the seed drive write will fail. □ **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.
Agents	Select the agents you want to replicate using the seed drive.

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.

10. Click **Start** to write the seed drive to the path you provided.

Replicating to a Core Managed by a Third Party

A third-party core is a target core that it managed and maintained by a 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.

This configuration applies to Hosted and Cloud Replication. The AppAssure 5 Core must be installed on all source core machines. If you are configuring AppAssure 5 for Multi-Point to Point replication, you must perform this task on all source cores.

Configuring Replication to a Target Core Managed by a Third Party

Complete the steps in this procedure to configure replication for a core that is managed by a third party.

To configure replication for a core managed by a third party

1. On the source core, navigate to the AppAssure 5 Core, and then click the Replication tab.
2. In the **Actions** drop-down menu, click **Add Remote Core**.
3. In the Select Replication Type dialog box, select the option, **I have a subscription to a third-party providing off-site backup and disaster recovery services, and wish to replicate my backups to that service**, and then enter the information as described in the following table.

Text Box	Description
Host Name	Enter the host name, IP address, or FQDN for the remote core machine.
Port	Enter the port number that was given to you by your third-party service provider. The default port number is 8006.

4. Click **Continue**.
5. In the Add Remote Core dialog box, do the following:
   a. Select agents to replicate.
   b. Select a repository for each agent.
   c. Enter your subscription email address and customer ID that was assigned to you by the service provider.
6. If you plan to perform the seeding process for the transfer of base data, select **Use a seed drive to perform initial transfer**.
7. Click **Submit Request**.

   NOTE: If you select Use a seed drive to perform initial transfer, the Copy to Seed Drive dialog box appears.

8. In the Copy to Seed Drive dialog box, enter the information for the seed drive as described in the table below.

Item	Description
Location	Enter the path to the drive on which you want to save the initial data, such as a local USB drive.
User Name	Enter the user name for connecting to the drive. **NOTE:** This is required if the seed drive is located on a network share.
Password	Enter the password for connecting to the drive. **NOTE:** This is required if the seed drive is located on a network share.
Maximum Size	Select one of the following options: □ The entire target □ A portion of the drive's available space. Then, to designate a portion of the drive, do the following: **a.** Enter the amount of space in the text box. **b.** Select the measurement.
Recycle action	In the event the path already contains a seed drive, select one of the following options: □ **Do not reuse**. Does not overwrite or clear any existing data from the location. If the location is not empty, the seed drive write will fail. □ **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.
Agents	Select agents you want to replicate using the seed drive.

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.

9. Click **Start** to write the seed drive to the path you provided.
10. Send the seed drive as directed by the third-party service provider.

Reviewing a Replication Request

After a user completes the procedure "Replicating to a Core Managed by a Third Party" on page 55, a replication request is sent from the source core to the third-party target core. As the third party, you can review the request, and then approve it to begin replication for your customer, or you can deny it to prevent replication from occurring.

Complete the following procedure to review a replication request on a third-party target core.

To deny a request without reviewing it, see "Ignoring a Replication Request" on page 58.

To review a replication request

1. On the target core, open the AppAssure 5 Core Console, and then click the Replication tab.
2. On the Replication tab, click **Pending Requests (#)**.

   The Pending Replication Requests section appears.
3. Under Pending Replication Requests, click the drop-down menu next to the request you want to review, and then click **Review**.

   The Review Replication Request window appears.

   NOTE: The information that appears in the Remote Core Identity section on this window is determined by the request completed by the customer.
4. On the Review Replication Request window, do one of the following:
   - To reject the request, click **Deny**.

     Replication is denied. Notification of denial appears under Alerts on the Events tab of the source core.
   - To approve the request, do the following:
     i. Confirm the Core Name, customer Email Address, and Customer ID, and then edit the information, if necessary.
     ii. Select the machines to which the approval applies, and then use the drop-down list to select the appropriate repository for each machine.
     iii. Optionally, in the Comment text box, enter any notes you want to display.
     iv. Click **Send Response**.

     Replication is accepted.

Ignoring a Replication Request

As a third-party service provider of a target core, you have the option to ignore a request for replication sent from a customer. This option could be used if a request was sent by mistake or if you want to reject a request without reviewing it.

For more information about replication requests, see "Reviewing a Replication Request" on page 57.

Complete the following procedure to ignore a replication request from a customer.

To ignore a replication request

1. On the target core, open the AppAssure 5 Core Console, and then click the Replication tab.
2. On the Replication tab, click **Pending Requests (#)**.

   The Pending Replication Requests section appears.

3. Under Pending Replication Requests, click the drop-down menu next to the request you want to ignore, and then click **Ignore**.
4. Notification that the request has been ignored is sent to the source core.

Consuming the Seed Drive on a Target Core

Complete the follow procedure to consume the data from the seed drive on the target core.


Note

This procedure is only necessary if a seed drive was created as part of "Configuring the Source Core to Replicate to a Self-Managed Target Core" on page 54 or "Replicating to a Core Managed by a Third Party" on page 55.

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. On the target core, open the AppAssure 5 Core Console, and then click the Replication tab.
3. On the Replication tab, under Incoming Replication, click the drop-down menu for the correct source core, and then click **Consume**.

   The Consume window appears.
4. Enter the information described in the following table.

Text Box	Description
Location	Enter the path to where the seed drive is located, such as a USB drive or network share. For example, D:\.
User name	Enter the user name for the shared drive or folder. **NOTE:** Required for network path only.
Password	Enter the password for the shared drive or folder. **NOTE:** Required for network path only.

5. 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, and it imports any comments made during Step 9 of "Replicating to a Self-Managed Core" on page 53.
6. On the Consume window, under Agent Names, select the machines for which you want to consume data.
7. Click **Consume**.

   To monitor the progress of consuming data, click 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. In this case, you may want to abandon the outstanding seed drive in favor different or more current seed data.

Complete the steps in the following procedure to abandon an outstanding seed drive.

This procedure removes the link to the outstanding drive from the AppAssure 5 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. On the source core, open the AppAssure 5 Core Console, and then click the Replication tab.
2. On the replication tab, click **Outstanding Seed Drives (#)**.

   The Outstanding seed drives section appears. It includes the name of the remote target core, the date and time on which the seed drive was created, and the data range of the recovery points included on the seed drive.
3. Under Outstanding seed drives, click the drop-down menu for the drive you want to abandon, and then click **Abandon**.

   The Outstanding Seed Drives window appears.
4. On the Outstanding Seed Drives window, click **Yes** to confirm.

   The seed drive is removed.

   If no more seed drives exist on the source core, the next time you visit the Replication tab, the Outstanding Seed Drives (#) link and Outstanding seed drives section are removed.

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 as described in the following table.

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 request is accepted by the MSP.	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 agent 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: □ **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 agents for replication, and create a seed drive for the new agent's initial transfer.
Incoming Replication	Lists all source machines from which the target receives replicated data. It includes the remote core name, state, machines, and progress.	On a target core, in the drop-down menu, you can select the following options: □ **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.

Pausing and Resuming Replication

You can pause replication temporarily for the source (outgoing) or target (incoming) cores.

To pause and resume replication

1. In the Core Console, click the Replication tab.
2. Under Outgoing Replication for a source core or Incoming Replication for a target core, click **Pause** to pause replication temporarily.
3. Click **Resume** to resume replication after it has been paused.

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 as described in the following table.

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 agents 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 agent 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" on page 63
- "Removing an Agent on the Target Core" on page 63

- "Removing a Target Core from Replication" on page 63
- "Removing a Source Core from Replication" on page 64

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

Removing an Agent from Replication on the Source Core

Complete the steps in this procedure to remove an agent from replication on the source core.

To remove an agent from replication on the source core

1. From the source core, open the AppAssure 5 Core Console, and click the Replication tab.
2. Expand the Outgoing Replication section.
3. In the drop-down menu for the agent machine that you want to remove from replication, click **Delete.**
4. In the Outgoing Replication dialog box, click **Yes** to confirm deletion.

Removing an Agent on the Target Core

Complete the steps in this procedure to remove an agent on the target core.

To remove an agent on the target core

1. On the target core, open the AppAssure 5 Core Console, and click the Replication tab.
2. Expand the Incoming Replication section.
3. In the drop-down menu for the agent machine that you want to remove from replication, click **Delete**, and then select one of the following options.

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

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

Removing a Target Core from Replication

Complete the steps in this procedure to remove a target core from replication.

To remove a target core from replication

1. On the source core, open the AppAssure 5 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

Complete the steps in this procedure to remove a source core from replication.

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 AppAssure 5 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. For more information, see "Mounting a Recovery Point for a Windows Machine" on page 136.
- Roll back to recovery points. For more information, see "Performing a Rollback" on page 153 or "Performing a Rollback for a Linux Machine by Using the Command Line" on page 154.

- Perform a virtual machine (VM) export. For more information, see "Exporting Backup Information for your Windows Machine to a Virtual Machine" on page 144.
- Perform a bare metal restore (BMR). For more information, see "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.
- Perform Failback (in the event you have a Failover/Failback replication environment set up). For more information, see "Performing Failback" on page 66.

Roadmap for Failover and Failback

When you encounter a disaster situation in which your source core and associated agent have failed, you can enable failover in AppAssure 5 to switch protection to your identical failover (target) core and launch a new (replicated) agent identical to the failed agent. Once your source core and agents have been repaired, you can then perform failback to restore the data from the failed-over core and agent back to the source core and agent. In AppAssure 5, failover and failback involve the following procedures.

- **Setting up your environment for failover**. See the section, "Setting Up an Environment for Failover" on page 65.
- **Perform failover for the target core and associated agent.** See the section, "Performing Failover on the Target Core" on page 66.
- **Restore a source core by performing failback.** See the section, "Performing Failover on the Target Core" on page 66

Setting Up an Environment for Failover

Setting up your environment for failover requires that you have a source and target AppAssure Core and associated agent set up for replication. Complete the steps in this procedure to set up replication for failover.

To set up an environment for failover

1. Install an AppAssure 5 Core for the source and install an AppAssure 5 Core for the target. For more information, see the AppAssure 5 Deployment Guide.
2. Install an AppAssure 5 Agent to be protected by the source core. For more information, see the AppAssure 5 Deployment Guide.
3. Create one repository on the source core and one repository on the target core. For more information, see "Creating a Repository" on page 37
4. Add the agent for protection under the source core.For more information, see "Protecting a Machine" on page 92.
5. Set up replication from the source to target core and replicate the protected agent with all recovery points. Follow the instructions in the section "Replicating to a Self-Managed Core" on page 53 to add the target core to which to replicate.

Performing Failover on the Target Core

When you encounter a disaster situation in which your source core and associated agents have failed, you can enable failover in AppAssure 5 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 AppAssure 5 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 status in this table for this machine changes to Failover.
4. Click the Machines tab, and then select the machine that has the associated AppAssure agent with recovery points.
5. Export the backup recovery point information on that agent to a virtual machine. For more information, see "Exporting Backup Information for your Windows Machine to a Virtual Machine" on page 144.
6. Shut down the machine that has the AppAssure agent.
7. Start the virtual machine that now includes the exported backup information. You need to wait for the device driver software to be installed.
8. Reboot the virtual machine and wait for the agent service to start.
9. Go back to the Core Console for the target core and verify that the new agent appears on the Machines tab under Protected Machines and on the Replication tab under Incoming Replication.
10. Force multiple snapshots, and verify they complete correctly. For more information, see "Forcing a Snapshot" on page 141.
11. You can now proceed with performing failback. See the following section, "Performing Failback" on page 66.

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 AppAssure 5 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 Failback Warnings dialog box opens to describe the steps you need to follow before you click the Start Failback button.

4. Click **Cancel**.

5. If the failed-over machine is running Microsoft SQL Server or Microsoft Exchange Server, stop those services.

6. In the Core Console for the target core, click the Tools tab.

7. Create an archive of the failed-over agent and output it to disk or a network share location. For more information, see the section, "Creating an Archive" on page 82.

8. After you create the archive, navigate to the Core Console on the newly repaired source core, and click the Tools tab.

9. Import the archive you just created in Step 7. For more information, see the section, "Importing an Archive" on page 82.

10. Go back to the Core Console on the target core, and click the Replication tab.

11. Under Incoming Replication, select the failover agent and expand the details.

12. On the Actions menu, click **Failback**.

13. In the Failback Warnings dialog box, click **Start Failback**.

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. For more information, see "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

    NOTE: When you launch the restore as described in, "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165, you will need to 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, see "Configuring Machine Settings" on page 98.

18. Navigate to the Core Console on the target core, and delete the agent from the Replication tab. For more information, see "Removing Replication" on page 62.

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. For more information, see the section, "Replicating to a Self-Managed Core" on page 53.

Managing Events

Managing core events assists with the monitoring of the health and usage of the AppAssure 5 Core. 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, email SMTP settings, repetition reduction, and event retention. The Notification Groups option in AppAssure 5 lets you 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

Complete the steps in this procedure to configure notification groups for events.

To configure notification groups

1. Navigate to the AppAssure 5 Core, and then select the Configuration tab.
2. From the Manage option, click **Events**.
3. Click **Add Group**.

   The Add Notification Group dialog box displays.

   The Add Notification Group dialog box contains a general description area and two tabs:

- Enable Events
- Notification Options

4. In the description area, enter basic information for the notification group, as described in the following table.

Text Box	Description
Name	Enter a name for event notification group. It is used to identify the event notification group.
Description	Enter a description for the event notification group. It is used to describe the purpose of the event notification group.

5. In the Enable Events tab, select the conditions for which event logs (alerts) to create and report. You can elect to create alerts for:
   - All Events
   - Appliance Events
   - BootCD
   - Security
   - DatabaseRetention
   - LocalMount
   - Clusters
   - Notification
   - Power Shell Scripting
   - Push Install
   - Nightly Jobs
   - Attachability
   - Jobs
   - Licensing
   - Log Truncation
   - Archive
   - Core Service
   - Export
   - Protection
   - Replication
   - Repository
   - Rollback
   - Rollup

6. In the Notification Options tab, specify how to handle the notification process.

   The following table describes the notification options.

Text Box	Description
Notify by email	Designate the recipients of the email notification. You can choose to specify separate multiple email addresses as well as blind and carbon copies. You can choose: □ To: □ CC: □ BCC:
Notify by Windows Event Log	Select this option if you want alerts to be reported through the Windows Event Log. It is used to specify whether the notification of alerts should be reported through the Windows Event Log.
Notify by sys logd	Select this option if you want alerts to be reported through syslogd. Specify the details for the syslogd in the following text boxes: □ Hostname: □ Port:
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.

7. Click **OK**.

Configuring an Email Server and Email Notification Template

Complete the steps in this procedure to configure an email server and email notification template.

You must also configure notification group settings, including enabling the **Notify by email** option, before email alert messages will be sent. For more information on specifying events to receive email alerts, see Configuring Notification Groups on page 68.

To configure an email server and email notification template

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Events**.

3. In the Email SMTP Settings pane, click **Change**.

   The Edit Email Notification Configuration dialog box appears.

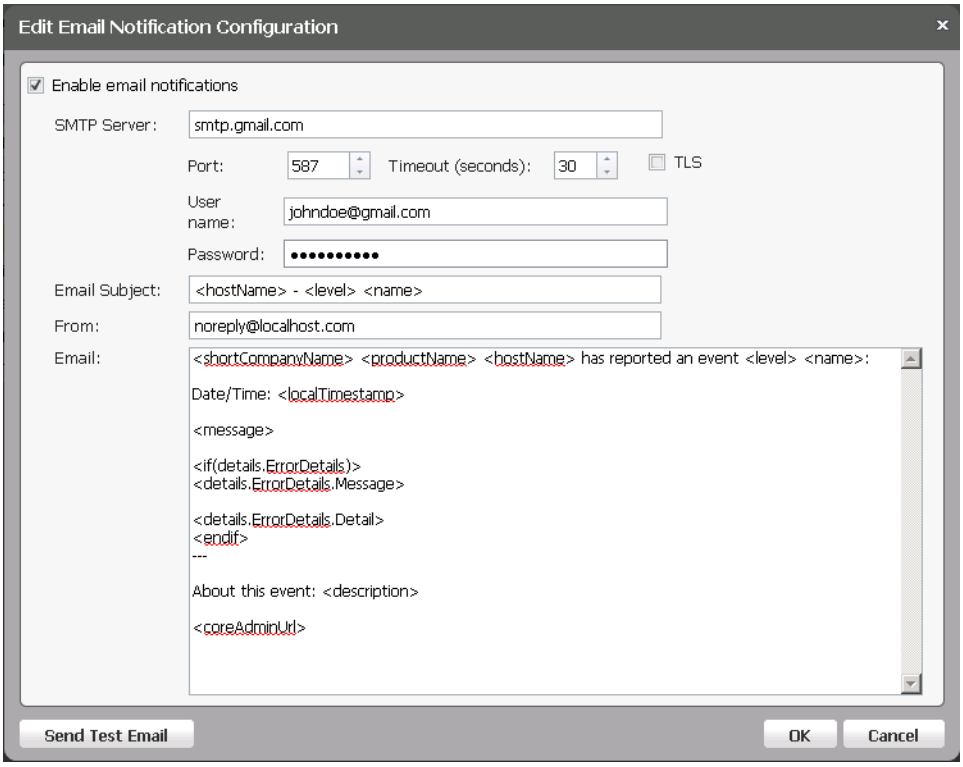

Figure 6. Edit Email Notification Configuration dialog box

4. Select **Enable Email Notifications**, and then enter details for the email server as described in the following table.

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.
User Name	Enter a user name for the email server.
Password	Enter the password associated with the user name required to access the email server.
Port	Enter a port number. It is used to identify the port for the email server; for example, the port 587 for Gmail. The default is 25.

Text Box	Description
Timeout (seconds)	Enter an integer value to specify how long to try to connect to the email server. It is used to establish the time in seconds before a timeout occurs. The default is 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).

5. Enter a subject for the email template. It is used to define the subject of the email notification template, for example, <hostname> - <level> <name>.
6. Enter the information for the body of the template which describes the event, when it occurred, and the severity.
7. Click **Send Test Email** and review the results.
8. Once you are satisfied with the results of the tests, click **OK**.

Configuring Repetition Reduction

Complete the steps in this procedure to configure repetition reduction for events.

To configure repetition reduction

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Events**.
3. From the Repetition Reduction area, click **Change**.

   The Repetition Reduction dialog box displays.
4. Select **Enable Repetition Reduction**.
5. In the Store events for X minutes text box, enter the number of minutes to store the events for repetition reduction.
6. Click **OK**.

Configuring Event Retention

Complete the steps in this procedure to configure retention for events.

To configure event retention

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Settings**.
3. Under Database Connection Settings, click **change**.

   The Database Connection Settings dialog box displays.

4. 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, 30 days (default).
5. Click **Save**.

Managing Recovery

The AppAssure 5 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 that all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot. Using application-aware snapshots in tandem with Recovery Assure 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

About System Information

AppAssure 5 lets you view information about the AppAssure 5 Core that includes system information, local and mounted volumes, and AppAssure engine connections.

If you want to dismount individual or all recovery points that are mounted locally on a core, you can accomplish this from the Mount option on the Tools tab. For more information about dismounting recovery points, see "Dismounting Select Recovery Points" on page 137 and "Dismounting All Recovery Points" on page 137.

Viewing System Information

Complete the steps in this procedure to view system information.

To view system information

1. Navigate to the AppAssure 5 Core, and then select the Tools tab.
2. From the Tools option, click **System Info**.

Downloading Installers

AppAssure 5 lets you download installers from the AppAssure 5 Core. From the Tools tab, you can choose to download the Agent Installer or the Local Mount Utility.

Note

For access to the Agent Installer, see Downloading the Agent Installer on page 75. For more information about deploying the Agent Installer, see the AppAssure 5 Deployment Guide. For access to the Local Mount Utility Installer, see About the Local Mount Utility on page 75 and for more information about the Local Mount Utility, see Downloading and Installing the Local Mount Utility on page 75.

About the Agent Installer

The Agent installer is used to install the AppAssure 5 Agent application on machines that are intended to be protected by the AppAssure 5 Core. If you determine that you have a machine that requires the Agent Installer, you can download the web installer from the Tools tab in the AppAssure 5 Core.

Note

Downloading of the Core is performed from the License Portal. For more information or to download the AppAssure 5 Core installer, visit https://licenseportal.com.

Downloading the Agent Installer

You can download and deploy the AppAssure 5 Agent Installer on any machine that will be protected by the AppAssure 5 Core. Complete the steps in this procedure to download the web installer.

To download the AppAssure 5 Agent installer

1. Download the AppAssure 5 Agent installer file from the AppAssure 5 License Portal or from the AppAssure 5 Core. For example:

   ```
   Agent-X64-5.2.1.xxxxx.exe
   ```

2. Click **Save File**.

   For more information about installing agents, see the AppAssure 5 Deployment Guide.

About the Local Mount Utility

The Local Mount Utility (LMU) is a downloadable application that lets you mount a recovery point on a remote AppAssure 5 Core from any machine. The light-weight utility includes the aavdisk and aavstor drivers, but it does not run as a service. When you install the utility, by default, it is installed in the directory C:\Program Files\AppRecovery\Local Mount Utility and a shortcut appears on the machine's desktop.

While the utility was designed for remote access to cores, you also can install the LMU on an AppAssure 5 Core. When it runs on a core, the application recognizes and displays all mounts from that core, including mounts performed through the AppAssure 5 Core Console. Likewise, mounts performed on the LMU also appear in the console.

Downloading and Installing the Local Mount Utility

AppAssure 5 lets you download the Local Mount Utility directly from the AppAssure 5 Core Console. Complete the following steps to download and install the utility.

To download and install the Local Mount Utility

1. From the machine on which you want to install the LMU, access the AppAssure 5 Core Console by entering the console URL into your browser and logging on with your user name and password.
2. From the AppAssure 5 Core Console, click the Tools tab.
3. From the Tools tab, click **Downloads**.
4. Under Local Mount Utility, click the **Download web installer** link.

5. From the Opening LocalMountUtility-Web.exe window, click **Save File.**

   The file saves to the local Downloads folder. In some browsers, the folder automatically opens.

6. From the Downloads folder, right-click the LocalMountUtility-Web executable and click **Open**.

   Depending on your machine's configuration, the User Account Control window could appear.

7. If the User Account Control window appears, click **Yes** to let the program make changes to the machine.

   The AppAssure Local Mount Utility Installation wizard launches.

8. On the AppAssure Local Mount Utility Installation wizard Welcome screen, click **Next** to continue to the License Agreement page.

9. On the License Agreement page, select **I accept the terms in the license agreement,** and then click **Next** to continue to the Prerequisites page.

10. On the Prerequisites page, install any necessary prerequisites and click **Next** to continue to the Installation Options page.

11. On the Installation Options page, complete the following tasks:

    a. Choose a destination folder for the LMU by clicking the **Change** button.

       NOTE: The default destination folder is C:\Program Files\AppRecovery\LocalMountUtility.

    b. Select whether or not to **Allow Local Mount Utility to automatically send diagnostic and usage information to AppAssure Software, Inc.**

    c. Click **Next** to continue to the Progress page and download the application.

    The application downloads to the destination folder, with progress displayed in the progress bar. When finished, the wizard automatically advances to the Completed page.

12. Click **Finish** to close the wizard.

Adding a Core to the Local Mount Utility

To mount a recovery point, you must add the Core to the LMU. There is no limit as to how many cores you can add.

Complete the following procedure to set up the LMU by adding a core.

To add a core to the Local Mount Utility

1. From the machine on which the LMU is installed, launch the LMU by double-clicking the desktop icon.

2. If the User Account Control window appears, click **Yes** to let the program to make changes to the machine.

3. In the upper-left corner of the AppAssure Local Mount Utility window, click **Add Core.**

4. In the Add Core dialog box, enter the requested credentials described in the following table.

Option	Description
Host name	The name of the Core from which you want to mount recovery points. **NOTE:** If installing the LMU on a core, the LMU automatically adds the localhost machine.
Port	The port number used to communicate with the Core. The default port number is 8006.
Use my Windows user credentials	Select this option if the credentials you use to access the Core are the same as your Windows credentials.
Use specific credentials	Select this option if the credentials you use to access the Core are different from your Windows credentials.
User name	The user name used to access the Core machine. **NOTE:** This option is only available if you choose to use specific credentials.
Password	The password used to access the Core machine. **NOTE:** This option is only available if you choose to use specific credentials.

5. Click **Connect.**

6. If adding multiple cores, repeat Step 3 through Step 5 as necessary.

Mounting a Recovery Point Using the Local Mount Utility

Before mounting a recovery point, the local mount utility (LMU) must connect to the Core on which the recovery point is stored. As described in the procedure "Adding a Core to the Local Mount Utility" on page 76, the number of cores that can be added to the LMU is unlimited; however, the application can connect to only one core at a time. For example, if you mount a recovery point of an agent protected by one core and then mount a recovery point of an agent protected by a different core, the LMU automatically disconnects from the first core to establish a connection with the second core.

Complete the following procedure to mount a recovery point on a remote core using the LMU.

To mount a recovery point using the Local Mount Utility

1. From the machine on which the LMU is installed, launch the LMU by double-clicking the desktop icon.

2. From the main AppAssure Local Mount Utility window, expand the Core in the navigation tree to reveal the protected agents.
3. From the navigation tree, select the agent from which you want to mount a recovery point.

   The recovery points displays in the main frame.
4. Expand the recovery point you want to mount to reveal individual disk volumes or databases.

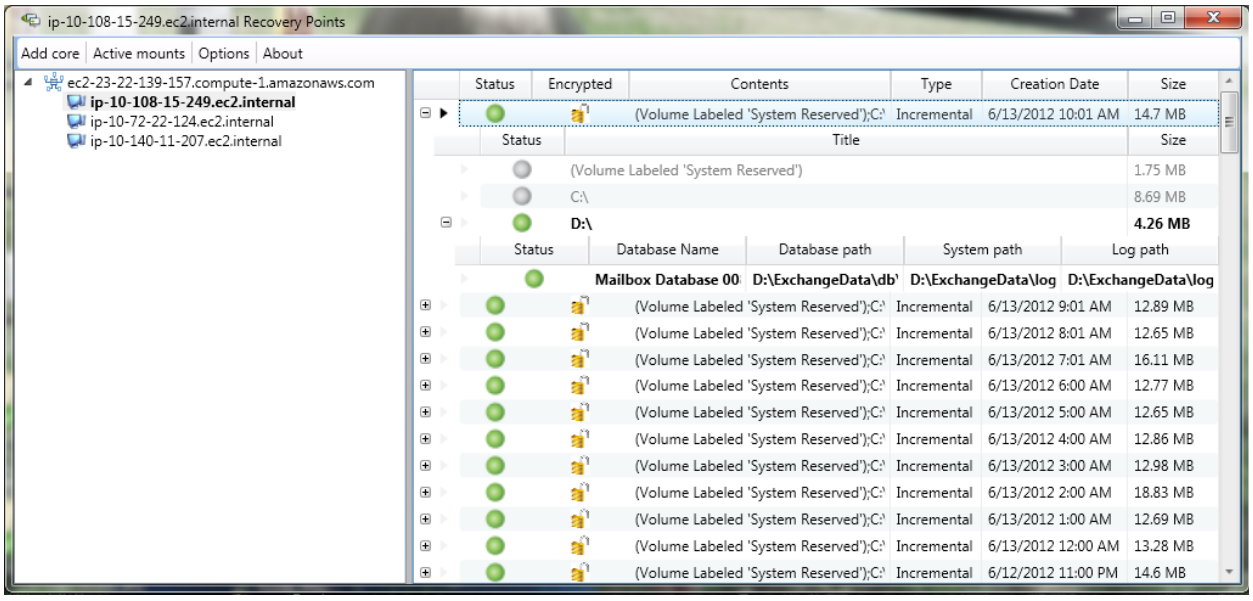

Figure 7. Local Mount Utility - recovery points

5. Right-click the recovery point you want to mount, and select one of the following options:
   - Mount
   - Mount writable
   - Mount with previous writes
   - Advanced mount
     i. If you selected Advanced Mount, then from the Advanced Mount window, complete the options described in the following table.

Option	Description
Mount point path	Click the **Browse** button to select a path for the recovery points other than the default mount point path.
Mount type	Select one of the following options: ☐ Mount read-only ☐ Mount writable ☐ Mount read-only with previous writes

ii. Click **Mount.**

The LMU automatically opens the folder containing the mounted recovery point.

NOTE: If you select a recovery point that is already mounted, the Mounting dialog will prompt whether to dismount the recovery point.

Exploring a Mounted Recovery Point Using the Local Mount Utility

Complete the following procedure to explore a recovery point that has remained mounted from a previous session.

This procedure is not necessary if you are exploring a recovery point immediately after mounting it, as the folder containing the recovery point automatically opens upon completion of the mounting procedure.

To explore a mounted recovery point using the Local Mount Utility

1. From the machine on which the LMU is installed, launch the LMU by double-clicking the desktop icon.
2. From the main Local Mount Recovery screen, click **Active mounts.**

   The Active Mounts window opens and displays all mounted recovery points.
3. Click **Explore** beside the recovery point from which you want to recover to open the folder of deduplicated volumes.

Dismounting a Recovery Point Using the Local Mount Utility

Complete the following procedure to dismount a recovery point on a remote core using the LMU.

To dismount a recovery point using the Local Mount Utility

1. From the machine on which the LMU is installed, double-click the Local Mount Utility desktop icon to launch the program.
2. From the main Local Mount Recovery screen, click **Active mounts.**

   The Active Mounts window opens and displays all mounted recovery points.
3. Do one of the following:
   - To dismount one recovery point, select a recovery point you want to dismount, and then click **Dismount**.
   - To dismount all mounted recovery points, click **Dismount all**, and then click **Yes** in the Dismount All dialog box to confirm.
4. To close the Active mounts window, click the X in the upper-right corner.

5. To minimize the LMU application, click the X in the upper-right corner of the Local Mount Utility window.
6. To close the LMU application, right-click the AppAssure Local Mount Utility icon in the LMU tray menu, and select **Exit**.

About the Local Mount Utility Tray Menu

The LMU tray menu is located in your desktop task bar. Right-click the icon to reveal the options described in the following table:

Option	Description
Browse Recovery Points	Opens the LMU main window.
Active Mounts	Opens the Active Mounts dialog box on top of the LMU main window.
Options	Opens the Options dialog box on top of the LMU main window. From the Options dialog box, you can change the Default mount point directory and the default Core credentials for the LMU user interface.
About	Reveals the Local Mount Utility licensing information.
Exit	Closes the LMU application. **NOTE:** Clicking the X in the upper corner of the main window of the LMU minimizes the application to the tray; it does not exit the application.

Using AppAssure 5 Core and Agent Options

By right-clicking the AppAssure 5 Core or agent in the main LMU screen, you can perform certain options. They include:

- Localhost Options
- Remote Core Options
- Agent Options

Accessing Localhost Options

Complete the step in this procedure to access Localhost options.

To access Localhost options

- Right-click the AppAssure 5 Core or agent and then click **Reconnect to core.**

  Information from the Core is updated and refreshed; for example, recently added agents.

Accessing Remote Core Options

Complete the steps in this procedure to access remote core options.

To access remote core options

- Right-click the AppAssure 5 Core or agent and then select one of the remote core options as described in the following table.

Option	Description
Reconnect to core	Refreshes and updates information from the Core, such as recently added agents.
Remove core	Deletes the Core from the Local Mount Utility.
Edit core	Opens the Edit Core window, where you can change the host name, port, and credentials.

Accessing Agent Options

Complete the steps in this procedure to access agent options.

To access agent options

- Right-click the AppAssure 5 Core or Agent, and then click **Refresh recovery points.**

  The list of recovery points for the selected agent updates.

Managing Retention Policies

Periodic backup snapshots of all the protected servers accumulate on the Core over time. The retention policies are used to retain backup snapshots for longer periods of time and to help with management of these backup snapshots. The retention policy is enforced by a nightly rollup process that helps in aging and deleting old backups. For information about configuring retention policies, see "Customizing Retention Policy Settings" on page 105.

About Archiving

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 non-compliance data. The archive feature in AppAssure 5 is used to support the extended retention for compliance and non-compliance data; and it is also used to seed replication data to a remote replica core.

Creating an Archive

Complete the steps in this procedure to create an archive.

To create an archive

1. Navigate to the AppAssure 5 Core, and then click the Tools tab.
2. From the Archive option, click **Create**.

   The Create Archive dialog box displays.
3. In the Create Archive dialog box, enter the details for the archive as described in the following table.

Text Box	Description
Date range	Select the to and from dates to specify the date range.
Location	Enter the location for the output. It is used to define the location path where you want the archive to reside. This can be a local disk or a network share; for example, d:\work\archive or \\servername\sharename for network paths. **NOTE:** If the output location is a network share, you will need to enter a user name and password for connecting to the share.
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.
Maximum Size	Enter how much space to use for the archive. You can select from: □ Entire Target, or □ Specific amount in MB, GB, or TB
Recycle action	Select the appropriate recycle action.
Comment	Enter any additional information that is necessary to capture for the archive. The comment will be displayed if you import the archive later.
Agent Names	Select one or more agents that you want to include in the archive.

4. Click **Archive**.

Importing an Archive

Complete the steps in this procedure to import an archive.

To import an archive

1. Navigate to the AppAssure 5 Core Console, and then select the Tools tab.
2. From the Archive option, click **Import**.

   The Import Archive dialog box displays.

3. In the Import Archive dialog box, enter the details for importing the archive as described in the following table.

Text Box	Description
Location	Select the location for importing the archive.
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.

4. Click **Check File** to validate the existence of the archive to import.

   The Restore dialog box displays.

5. In the Core list, verify the name of the source core.
6. Select the agents to be imported from the archive.
7. Select the repository.
8. Click **Restore** to import the archive.

Managing SQL Attachability

The SQL attachability configuration enables the AppAssure 5 Core to attach SQL database and log files in a snapshot of a SQL server using a local instance of Microsoft SQL Server. The attachability test lets the Core check for the consistency of the SQL databases and ensures that all data files (MDF and LDF files) are available in the backup snapshot. Attachability checks can be run on demand for specific recovery points or as part of a nightly job.

Attachability requires a local instance of Microsoft SQL Server on the AppAssure Core machine. This instance must be a fully licensed version of SQL Server procured from Microsoft or through a licensed reseller. Microsoft does not allow the use of passive SQL licenses.

Attachability supports SQL Server 2005, 2008, 2008 R2, and 2012. The account used to perform the test must be granted the sysadmin role on the SQL Server instance.

The SQL Server on-disk storage format is the same in both 64-bit and 32-bit environments and attachability works across both versions. A database that is detached from a server instance that is running in one environment can be attached on a server instance that runs in another environment.

The version of SQL Server on the Core must be equal to or newer than the SQL Server version on all of the agent machines with SQL Server installed.

Configuring SQL Attachability Settings

Prior to running attachability checks on protected SQL databases, you must first select a local instance of SQL Server on the Core machine that will be used to perform the checks against the agent machine.

Attachability requires a local instance of Microsoft SQL Server on the AppAssure Core machine. This instance must be a fully licensed version of SQL Server procured from Microsoft or through a licensed reseller. Microsoft does not allow the use of passive SQL licenses.

Complete the steps in this procedure to configure SQL attachability settings.

To configure SQL attachability settings

1. Navigate to the AppAssure 5 Core, and then click the Configuration tab.
2. From the Manage option, click **Attachability**.

   The Attachability Check Settings window displays.
3. Select the local SQL Server instance to use to perform attachability checks for the protected SQL databases.

   You can choose from:

   - SQL Server 2005
   - SQL Server 2008
   - SQL Server 2008 R2
   - SQL Server 2012
4. Select the credential type. You can select from:

   - Windows, or
   - SQL
5. Specify the credentials with administrative privileges for the Windows or SQL Server instances, as described in the following table.

Text Box	Description
User Name	Enter a user name for logon permissions to the SQL server.
Password	Enter a password for SQL attachability. It is used to control logon activity.

6. Click **Test Connection**.

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

   Attachability checks are now available to be run on the protected SQL Server databases.

Configuring Nightly SQL Attachability Checks and Log Truncation

Complete the steps in this procedure to have the system perform nightly attachability checks for the SQL Server recovery points.

To configure nightly SQL attachability checks and log truncation

1. In the left navigation area of the AppAssure 5 Core, select the machine for which you want to configure nightly attachability checks and log truncation.

   The Summary tab for the selected machine displays.

2. Click **SQL Server Settings**.

3. Select or clear the following SQL Server settings based on the needs of your organization:

   - Enable nightly attachability check
   - Enable nightly log truncation

4. Click **OK**.

   The attachability and log truncation settings take effect for the protected SQL Server.

   NOTE: These steps need to be performed for each of the protected machines under the Core. For more information on forcing log truncation, see "Forcing Log Truncation" on page 88.

Forcing a SQL Server Attachability Check

Complete the steps in this procedure to force the system to perform an attachability check for a specific SQL server recovery point.

To force a SQL Server attachability check

1. In the left navigation area of the AppAssure 5 Core, select the machine for which you want to force the attachability check and click the **Recovery Points** tab.

2. Click the right angle bracket **>** symbol next to a recovery point in the list to expand the view.

3. Click **Force Attachability Check.**

   The Force Attachability Check window appears for you to indicate that you want to force an attachability check.

4. Click **Yes**.

   The system performs the attachability check.

   NOTE: For information on how to view the status of the attachability checks, see "Viewing Events and Alerts" on page 186.

Managing Exchange Database Mountability Checks and Log Truncation

When using AppAssure 5 to back up Microsoft Exchange Servers, mountability checks can be performed on all Exchange databases after every snapshot. This corruption detection feature alerts administrators of potential failures and ensures that all data on the Exchange servers will be recovered successfully in the event of a failure.

The mountability checks and log truncation features only apply to Microsoft Exchange 2007, 2010, and 2013. Additionally, the AppAssure 5 Agent service account must be assigned the Organizational Administrator role in Exchange.

Configuring Exchange Database Mountability and Log Truncation

You can view, enable, or disable Exchange database server settings, including automatic mountability check, nightly checksum check, or nightly log truncation.

Complete the steps in this procedure to configure settings for Exchange database mountability and log truncation.

To configure Exchange database mountability and log truncation

1. In the left navigation area of the AppAssure 5 Core, select the machine for which you want to configure mountability checks and log truncation.

   The Summary tab for the selected machine displays.

2. Click **Exchange Server Settings**.

   The Exchange Server Settings dialog box displays

3. Select or clear the following Exchange Server settings based on the needs of your organization:

   - Enable automatic mountability check
   - Enable nightly checksum check
   - Enable nightly log truncation

4. Click **OK**.

   The mountability and log truncation settings take effect for the protected Exchange server.

   NOTE: For information on forcing log truncation, see "Forcing Log Truncation" on page 88.

Forcing a Mountability Check

Complete the steps in this procedure to force the system to perform a mountability check for a specific Exchange server recovery point.

To force a mountability check

1. In the left navigation area of the AppAssure Core Console, select the machine for which you want to force the mountability check, and then click the **Recovery Points** tab.
2. Click the right angle bracket **>** symbol next to a recovery point in the list to expand the view.
3. Click **Force Mountability Check**.

   A pop-up window displays asking if you want to force a mountability check.
4. Click **Yes**.

   The system performs the mountability check.

   NOTE: For instructions on how to view the status of the attachability checks, see "Viewing Events and Alerts" on page 186.

Forcing Checksum Checks

Complete the steps in this procedure to force the system to perform a checksum check for a specific Exchange server recovery point.

To force a checksum check

1. In the left navigation area of the AppAssure Core Console, select the machine for which you want to force the checksum check, and then click the **Recovery Points** tab.
2. Click the right angle bracket **>** symbol next to a recovery point in the list to expand the view.
3. Click **Force Checksum Check**.

   The Force Attachability Check window appears for you to indicate that you want to force a checksum check.
4. Click **Yes**.

   The system performs the checksum check.

   NOTE: For information on how to view the status of the attachability checks, see Viewing Events and Alerts on page 186.

Forcing Log Truncation

Complete the steps in this procedure to force log truncation.

This option is only available for Exchange or SQL machines.

To force log truncation

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to truncate the log.
   - Or, in the navigation pane, select the machine you want to truncate the log.
3. In the **Actions** drop-down menu for that machine, click **Force Log Truncation**.
4. Confirm whether to proceed with forcing log truncation.

Recovery Point Status Indicators

Once a recovery point is created on a protected SQL or Exchange server, the application displays a corresponding color status indicator in the Recovery Points table. The color that displays is based on the check settings for the protected machine and the success or failure of those checks, as described in the following Recovery Status Point Colors for SQL Databases and Recovery Status Point colors for Exchange Database tables.

For more information on viewing Recovery Points, see Viewing Recovery Points on page 133.

Recovery Status Point Colors for SQL Databases

The following table lists the status indicators that display for SQL databases.

Status Color	Description
White	Indicates that one of the following conditions exist: □ An SQL database did not exist □ Attachability checks were not enabled, or □ Attachability checks have not yet been run.
Yellow	Indicates that the SQL database was offline and a check was not possible.
Red	Indicates that the attachability check failed.
Green	Indicates that the attachability check passed.

Recovery Status Point Colors for Exchange Databases

The following table lists the status indicators that display for Exchange databases.

Status Color	Description
White	Indicates that one of the following conditions exist: □ An Exchange database did not exist, or □ Mountability checks were not enabled. **NOTE:** This can apply to certain volumes within a recovery point.
Yellow	Indicates that the Exchange database mountability checks are enabled, but the checks have not yet been run.
Red	Indicates that either the mountability or checksum checks failed on at least one database.
Green	Indicates that the mountability check passed or that the checksum check passed.

Recovery points that do not have an Exchange or SQL database associated with it will appear with a white status indicator. In situations where both an Exchange and SQL database exists for the recovery point, the most severe status indicator displays for the recovery point.

This page intentionally left blank.

4 Protecting Workstations and Servers

This chapter describes how to protect, configure, and manage the agent machines in your AppAssure environment. It includes the following sections:

About Protecting Workstations and Servers

To protect your data using AppAssure 5, you need to add the workstations and servers for protection in the AppAssure 5 Core Console; for example, your Exchange server, SQL Server, your Linux server, and so on.

Note

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

In the AppAssure 5 Core Console, you can identify the machine on which an AppAssure 5 Agent is installed and specify which volumes, for example, a Microsoft Windows Storage Space, to protect. You can define the schedules for protection, add additional security measures such as encryption, and much more. For more information on how to access the AppAssure 5 Core Console to protect workstations and servers, see "Protecting a Machine" on page 92.

Note

Windows 8, 8.1 and Windows Server 2012, 2012 R2 operating systems that are booted from FAT32 EFI partitions are not available for protection or recovery, nor are Resilient File System (ReFS) volumes. Bare metal restore of Storage Spaces disks configuration (a feature of Windows 8.1) is also not supported in this release. For details, see the AppAssure 5 Deployment Guide.

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 5 Agent software installed in order to be protected. You can choose to install the Agent software prior to this procedure, or you can deploy the software to the agent as you define protection in the Connection dialog box. For specific steps to install the agent software during the process of protecting a machine, see "Deploying the Agent Software When Protecting an Agent" on page 94.

When you add protection, you need to 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" on page 129.

To protect a machine

1. If you did not do so after installation, reboot the machine on which the AppAssure 5 Agent software is installed.
2. On the core machine, navigate to the AppAssure 5 Core Console, and do one of the following:
   - From the Home tab, under Protected machines, click **Protect Machine**.

- Select the Machines tab, and in the **Actions** drop-down menu, click **Protect Machine.**

The Connect dialog box displays.

3. In the Connect dialog box, 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 5 Core communicates with the Agent on the machine. The default port number is 8006.
User name	The user name used to connect to this machine; for example, administrator.
Password	The password used to connect to this machine.

4. Click **Connect** to connect to this machine.
5. If the Agent software is not yet on installed on the machine that you designated, pause here and follow the procedure "Deploying the Agent Software When Protecting an Agent" on page 94. Ensure that you restart the agent machine after deploying the Agent software. Then resume with the next step.
6. In the Protect dialog box, edit the settings as needed, as described in the following table.

Text Box	Description
Display Name	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 AppAssure 5 Core Console. **NOTE:** You can also change the display name later by accessing the Configuration tab for an existing machine.
Repository	Select the repository on the AppAssure 5 Core in which to store the data from this machine.
Encryption Key	Specify whether encryption should be applied to the data for every volume on this machine to be stored in the repository. **NOTE:** The encryption settings for a repository are defined under the Configuration tab in the AppAssure 5 Core Console.

Text Box	Description
Initially pause protection	Once you add a machine for protection, AppAssure 5 automatically begins the process of taking a base snapshot of data. You can select this check box to pause protection initially. You then need to force a snapshot manually when you are ready to start protecting your data. For more information, see "Forcing a Snapshot" on page 141.
Volume Groups	Under Volume Groups, you can define which volumes you want to protect, and establish a protection schedule. To set a default protection schedule of every 60 minutes for all volumes on the machine, click **Apply Default**. You can also select any volume on the machine and define protection parameters for it individually. Initial settings apply a default protection schedule of every 60 minutes. To modify the schedule for any volume, click **Edit** for that volume. You can then further define the interval between snapshots (including defining a separate schedule for weekends) or specify a daily time to begin a snapshot. For more information on editing the protection schedule for a selected volume, see "Creating Custom Schedules for Volumes" on page 96.

7. Click **Protect**.

   The first time protection is added for a machine, a base image (that is, a snapshot of all the data in the protected volumes) immediately begins to transfer to the repository on the AppAssure 5 Core, unless you specified to initially pause protection.

If you protected a Linux machine, note that you should not unmount a protected volume by hand. In the event you need to do this, 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 would need to be in a form similar to this example: `/dev/sda1`.

Deploying the Agent Software When Protecting an Agent

You can download and deploy agents when you add an agent for protection.

This step is not required if you have already installed the Agent software on a machine you want to protect.

Complete steps in this procedure to deploy the AppAssure 5 Agent software on the selected machine when adding a machine to protection.

To deploy agents when protecting an agent.

1. From the Connect dialog box, after entering the appropriate connection settings, click **Connect**.

   The Deploy Agent dialog box appears.

2. Click **Yes** to deploy the Agent software remotely to the selected machine.

   The Deploy Agent dialog box appears.

3. In the Deploy Agent dialog box, enter logon and protection settings as defined in the following table.

Text Box	Description
Host name	The host name or IP address of the machine that you want to protect.
Port	The port number on which the AppAssure 5 Core communicates with the Agent on the machine. The default port number is 8006.
User name	The user name used to connect to this machine; for example, administrator.
Password	The password used to connect to this machine.
Display Name	Enter a new name for the machine to be displayed in the AppAssure 5 Core Console. This could be the same value as the host name. **NOTE:** You can also change the display name later by accessing the Configuration tab for an existing machine in the AppAssure 5 Core Console.
Protect machine after install	If this option is selected, once you add a machine for protection, AppAssure 5 automatically begins the process of taking a base snapshot of data. If you clear this option, then you will need to force a snapshot manually when you are ready to start protecting your data. For more information, see "Forcing a Snapshot" on page 141.
Repository	Select the repository on the AppAssure 5 Core in which to store the data from this machine. **NOTE:** You can store the data from multiple agents in a single repository. This allows you to deduplicate among multiple protected machines.
Encryption key	Specify whether encryption should be applied to the data for every volume on this machine to be stored in the repository. **NOTE:** The encryption settings for a repository are defined under the Configuration tab in the AppAssure 5 Core Console.

4. Click **Deploy**.

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

You must reboot the agent machine after installing the Agent software. This ensures that the drivers load properly, allowing the AppAssure 5 Core to monitor file activity on the agent and ensure successful backups. If you cannot reboot at the time of installation, schedule a reboot later in the day to ensure AppAssure 5 is able to fully protect your data.

Creating Custom Schedules for Volumes

Complete the steps in this procedure to create custom schedules for volumes on a machine.

To create custom schedules for volumes

1. In the Protect Machine dialog box (see the section, "Protecting a Machine" on page 92, for information about accessing this dialog box), under Volume Groups, select a volume for protection, and then click **Edit**.

   The Protection Schedule dialog box displays.

2. In the Protection Schedule dialog box, select one of the following schedule options for protecting your data as described in the following table.

Text Box	Description
Interval	You can choose from: □ **Weekday**. To protect data on a specific interval, select **Interval**, and then: □ To customize when to protect data during peak times, you can specify a Start Time, End Time, and an Interval from the drop-down menus. □ To protect data during off-peak times, select **Protection interval during off-peak times**, and then select an interval for protection from the **Time** drop-down menu. □ **Weekends**. To protect data during weekends as well, select **Protection interval during weekends**, and then select an interval from the drop-down menu.
Daily	To protect data on a daily basis, select the **Daily protection** option, and then, in the Time drop-down menu, select a time to start protecting data.
No Protection	To remove protection from this volume, select the **No Protection** option.

   If you want to apply these custom settings to all the volumes on this machine, select **Apply to All Volumes**.

3. When you have made all necessary changes, click **OK**.
4. Repeat Step 2 and Step 3 for any additional volumes you want to customize.
5. In the Protect Machine dialog box, click **Protect**.

Modifying Exchange Server Settings

If you are protecting data from a Microsoft Exchange server, you need to configure additional settings in the AppAssure 5 Core Console.

To modify Exchange server settings

1. Once you have added the Exchange Server machine for protection, navigate to the AppAssure 5 Core Console and select the machine in the Navigation pane.

   The Summary tab displays for the machine.

2. From the Summary tab, click the **Exchange Server Settings** link.

   The Exchange Server Settings dialog box appears.

3. In the Exchange Server Settings dialog box, you can select or clear the following settings:

   - Enable automatic mountability check, or
   - Enable nightly checksum check. You can further customize this setting by selecting the following:
     - Automatically truncate Exchange logs after successful checksum check
     - Truncate log before checksum check completes

4. You can also modify the logon credentials for your Exchange Server. To do so, scroll down to the Exchange Server Information section and then click **Change Credentials**.

   The Set Exchange Credentials dialog box appears.

5. Enter your new credentials and then click **OK**.

Modifying SQL Server Settings

If you are protecting data from Microsoft SQL Server, there are additional settings you need to configure in the AppAssure 5 Core Console.

To modify SQL server settings

1. Once you have added the SQL Server machine for protection, from the AppAssure 5 Core Console, select the machine in the Navigation pane.

   The Summary tab displays for the machine.

2. From the Summary tab, click the **SQL Server settings** link.

   The SQL Server Settings dialog box appears.

3. In the SQL Server Settings dialog box, edit the following settings, as needed:

   - Enable nightly attachability check
   - Truncate log after successful attachability check (simple recovery model only)

4. You can also modify the logon credentials for SQL Server. To do so, scroll down to the SQL Server Information table and then click **Change Credentials**.

   The Set SQL Server Credentials dialog box appears.

5. Enter your new credentials, and then click **OK**.

Pausing and Resuming Protection

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

To pause and resume protection

1. In the AppAssure 5 Core Console, click the Machines tab.
2. Select the machine for which you want to pause protection.

   The Summary tab for this machine displays.

3. In the **Actions** drop-down menu for that machine, click **Pause.**
4. To resume protection, click **Resume** in the Actions menu.

Configuring Machine Settings

Once you have added protection for machines in AppAssure, you can easily modify basic machine configuration settings (name, host name, and so on), protection settings (changing the protection schedule for volumes on the machine, adding or removing volumes, or pausing protection), and more. This section describes the various ways you can view and modify machine settings in AppAssure.

Viewing and Modifying Configuration Settings

Complete the steps in this procedure to modify and view configuration settings.

This task is also a step in the "Process of Modifying Cluster Node Settings" on page 193.

To view and modify configuration settings

1. Once you have added a protected machine, perform one of the following:
   a. From the AppAssure 5 Core Console, click the Machines tab and then click the hyperlink for the machine you want to modify.
   b. Or, from the Navigation pane, select the machine you want to modify.
2. Click the Configuration tab.

   The Settings page displays.

3. Click **Change** to modify the machine settings as described in the following table.

Text Box	Description
Display Name	Enter a display name for the machine. A name for this machine to be displayed in the AppAssure 5 Core Console. By default, this is the host name of the machine. You can change this to something more user-friendly if needed.
Host Name	Enter a host name for the machine.
Port	Enter a port number for the machine. The port is used by the Core to communicate with this machine.
Encryption Key	Edit the encryption key if necessary. Specifies whether encryption should be applied to the data for every volume on the machine that is stored in the repository.
Repository	Select a repository for the recovery points. Displays the repository on the AppAssure 5 Core in which to store the data from this machine. **NOTE:** This setting can only be changed if there are no recovery points or the previous repository is missing.

Viewing System Information for a Machine

The AppAssure 5 Core Console provides you with an at-a-glance view of all of the machines that are being protected by including a list of the machines as well as each machine's status.

Complete the steps in this procedure to view system information for a protected machine.

To view system information for a machine

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to view.
   - Or, in the navigation pane, select the machine you want to view.
3. Click the Tools tab, and then click **System Info**.

   The information about the machine displays in the System Information page. The details that display include the following:
   - Host Name
   - OS Version
   - OS Architecture

- Memory (Physical)
- Display Name
- Fully Qualified Domain Name

Detailed information about the volumes contained on this machine also displays and includes:

- Processors
- Type of Processors
- Network Adapters
- IP Addresses associated with this machine

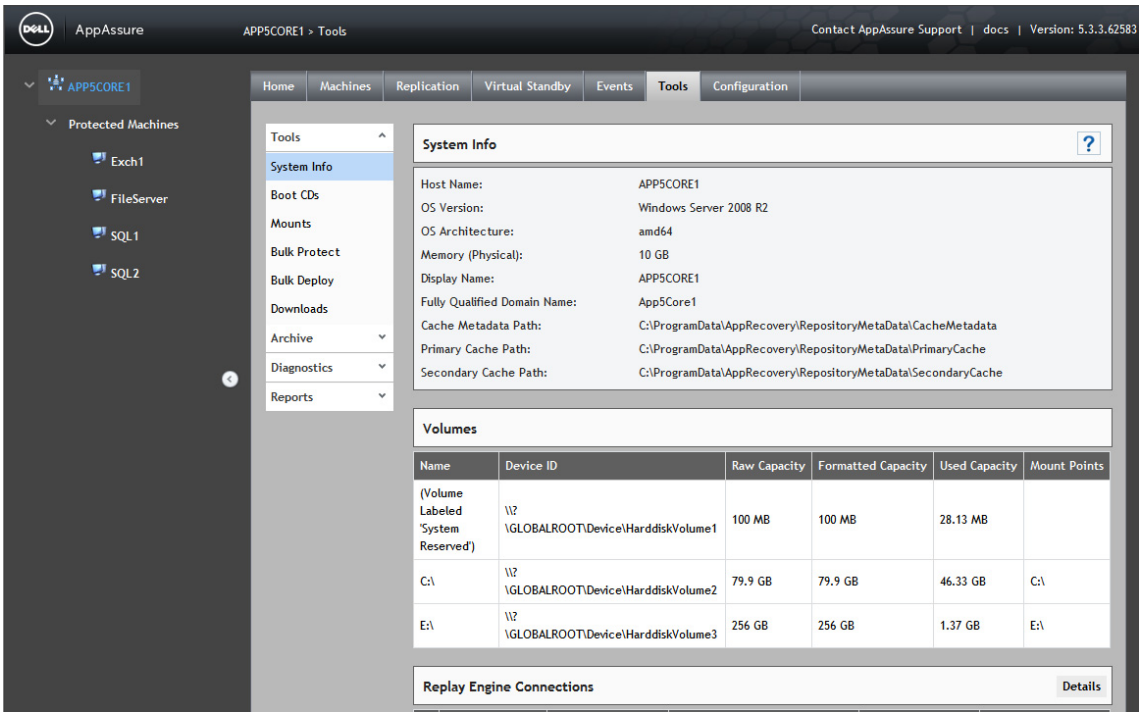

Figure 8. System information

Configuring Notification Groups for System Events

In AppAssure 5, you can configure how system events are reported for your machine by creating notification groups. These events could be system alerts, errors, and so on.

To configure notification groups for system events

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.

2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to modify.
   - Or, in the navigation pane, select the machine you want to modify.

   The Summary Tab appears.
3. Click the Configuration tab, and then click **Events.**

   The Notification Groups page displays.
4. Click **Use custom alert settings** and then click **Apply**.

   The Custom Notification Groups screen appears.
5. Click **Add Group** to add new notification groups for sending a list of system events.

   The Add Notification Group dialog box displays.

   NOTE: To use the default alert settings, select the **Use Core alert settings** option.
6. Add the notification options as described in the following table.

Text Box	Description
Name	Enter a name for the notification group.
Description	Enter a description for the notification group.

Text Box	Description
Enable Events	Select which events to share with this notification group. You can select **All** or select a subset of events to include: □ BootCd □ LocalMount □ Metadata □ Clusters □ Notification □ PowerShellScripting □ PushInstall □ Attachability □ Jobs □ Licensing □ LogTruncation □ Archive □ CoreService □ Export □ Protection □ Replication □ Rollback □ Rollup You can also choose to select by type: □ Info □ Warning □ Error **NOTE:** When you choose to select by type, by default, the appropriate events are automatically enabled. For example, if you choose Warning, the Attachability, Jobs, Licensing, Archive, CoreService, Export, Protection, Replication, and Rollback events are enabled.
Notification Options	Select the method to specify how to handle notifications. You can choose from the following options: □ **Notify by Email**. You would need to specify to which email addresses to send the events in the To, CC and, optionally, BCC text boxes. **NOTE:** To receive mail, SMTP must be previously configured. □ **Notify by Windows Event log**. The Windows Event log controls the notification. □ **Notify by syslogd**. You would need to specify to which host name and port to send the events. □ **Host**. Enter the host name for the server. □ **Port**. Enter a port number for communicating with the server. □ **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.

7. Click **OK** to save your changes.
8. To edit an existing notification group, click **Edit** next to the notification group that you want to edit.

   The Edit Notification Group dialog box displays for you to edit the settings.

Editing Notification Groups for System Events

Complete the steps in the following procedure to edit notification groups for system events.

To edit notification groups for system events

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   a. Click the hyperlink for the machine you want to modify.
   b. Or, in the navigation pane, select the machine you want to modify.

   The Summary Tab appears.
3. Click the Configuration tab, and then click **Events.**
4. Click **Use custom alert settings** and then click **Apply**.

   The Custom Notification Groups screen appears.
5. Click the Edit icon under the Action column.

   The Edit Notification Group dialog box displays.
6. Edit the notification options as described in the following table.

Text Box	Description
Name	Represents the name of the notification group. **NOTE:** You cannot edit the name of the notification group.
Description	Enter a description for the notification group.

Text Box	Description
Enable Events	Select which events to share with the notification group. You can select **All** or select a subset of events to include: □ BootCd □ LocalMount □ Metadata □ Clusters □ Notification □ PowerShellScripting □ PushInstall □ Attachability □ Jobs □ Licensing □ LogTruncation □ Archive □ CoreService □ Export □ Protection □ Replication □ Repository □ Rollback □ Rollup You can also choose to select by type, which are: Info, Warning, and Error. **NOTE:** When you choose to select by type, by default, the appropriate events are automatically enabled. For example, if you choose Warning, the Attachability, Jobs, Licensing, Archive, CoreService, Export, Protection, Replication, and Rollback events are enabled.
Notification Options	Select the method to specify how to handle notifications. You can choose from the following options: □ **Notify by Email**. You would need to specify the email addresses to which to send the events in the To, CC and optionally, BCC text boxes. **NOTE:** To receive email, SMTP must be previously configured. □ **Notify by Windows Event log**. The Windows Event log controls the notification. □ **Notify by syslogd**. You would need to specify the host name and port to which to send the events. □ **Host**. Enter the host name for the server. □ **Port**. Enter a port number for communicating with the server. □ **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.

7. Click **OK**.

Customizing Retention Policy Settings

The retention policy for a machine specifies how long the recovery points for an agent machine are stored in the repository. Retention policies are used to retain backup snapshots for longer periods of time and to help manage these backup snapshots. The retention policy is enforced by a rollup process that helps with aging and deleting old backups.

This task is also a step in the "Process of Modifying Cluster Node Settings" on page 193.

To customize retention policy settings

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   a. Click the hyperlink for the machine you want to modify.
   b. Or, in the navigation pane, select the machine you want to modify.

   The Summary Tab appears.
3. Click the Configuration tab, and then click **Retention Policy**.

   NOTE: To use the default retention policy configured for the Core, make sure the **Use Core default retention policy** option is selected.

   The Retention Policy screen appears.
4. Select **Use custom retention policy** to set the customized policies.

   The Custom Retention Policy screen expands to display the retention policy options.

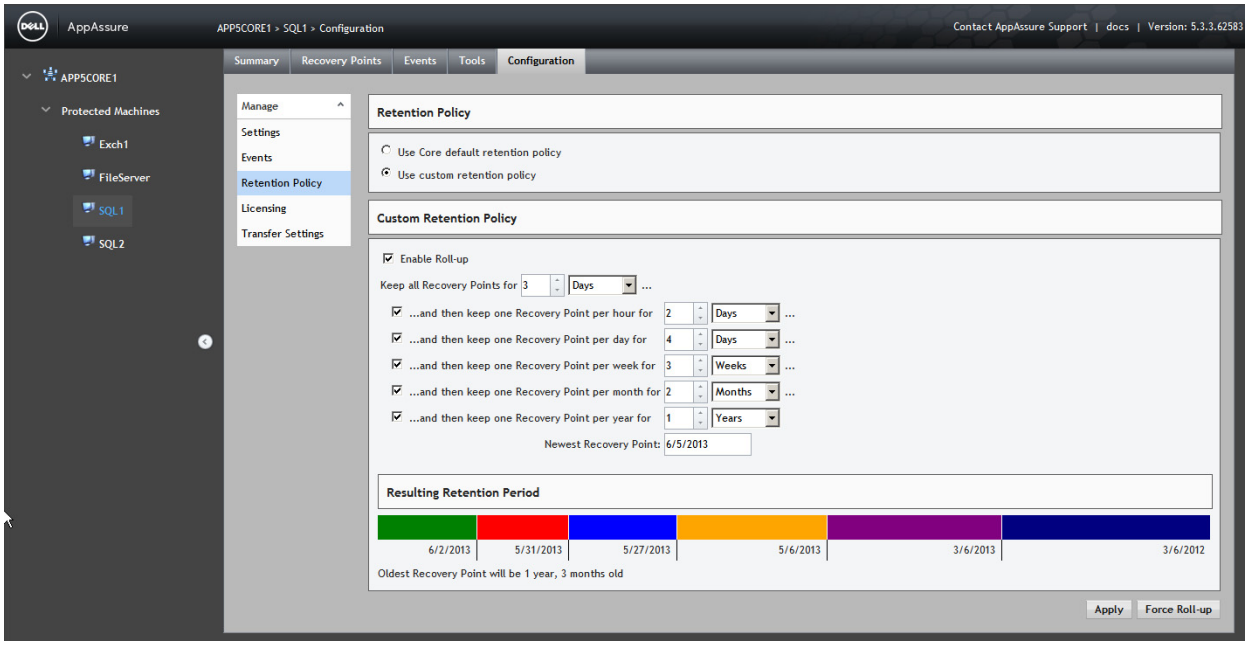

Figure 9. Custom Retention Policy

5. Select **Enable Rollup**, and specify the time intervals for retaining the backup data as needed. The retention policy options are described in the following table.

Text Box	Description
Keep all Recovery Points for n [retention time period]	Specifies the retention period for the recovery points. Enter a number to represent the retention period and then select the time period. The default is 3. You can choose from: □ Days □ Weeks □ Months □ Years
...and then keep one Recovery Point per hour for n [retention time period]	Provides a more granular level of retention. It is used as a building block with the primary setting to further define how long recovery points are maintained. Enter a number to represent the retention period and then select the time period. The default is 2. You can choose from: □ Days □ Weeks □ Months □ Years
...and then keep one Recovery Point per day for n [retention time period]	Provides a more granular level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number to represent the retention period and then select the time period. The default is 4. You can choose from: □ Days □ Weeks □ Months □ Years
...and then keep one Recovery Point per week for n [retention time period]	Provides a more granular level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number to represent the retention period and then select the time period. The default is 3. You can choose from: □ Weeks □ Months □ Years

Text Box	Description
...and then keep one Recovery Point per month for n [retention time period]	Provides a more granular level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number to represent the retention period and then select the time period. The default is 2. You can choose from: □ Months □ Years
...and then keep one Recovery Point per year for n [retention time period]	Enter a number to represent the retention period and then select the time period. You can choose from: □ Years

The Newest Recovery Point text box displays the most recent recovery point. The oldest recovery point would be determined by the retention policy settings.

The following is an example of how the retention period is calculated.

Keep all recovery points for 3 days.

...and then keep one recovery point per hour for 3 days

...and then keep one recovery point per day for 4 days

...and then keep one recovery point per week for 3 weeks

...and then keep one recovery point per month for 2 months

...and then keep one recovery point per month for 1 year

Newest Recovery Point is set to the current day, month, and year.

In this example, the oldest recovery point would be 1 year, 4 months, and 6 days old.

6. Click **Apply** to save your changes.
7. Click **Force Rollup** to perform a rollup based on the current retention policy for the machine, or let the retention policy you defined be applied during the nightly rollup.

Viewing License Information

You can view current license status information for the AppAssure 5 Agent software installed on a machine.

To view license information

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to view.
   - Or, in the navigation pane, select the machine you want to view.

3. Click the Configuration tab, and then click **Licensing**.

   The Status screen appears and presents the details about the product licensing.

Modifying Protection Schedules

In AppAssure 5, you can modify the protection schedules for specific volumes on a machine.

To modify protection schedules

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to modify.
   - Or, in the navigation pane, select the machine you want to modify.
3. In the Volumes table on the Summary tab for the machine, click the hyperlink for the protection schedule for the volume you want to customize.

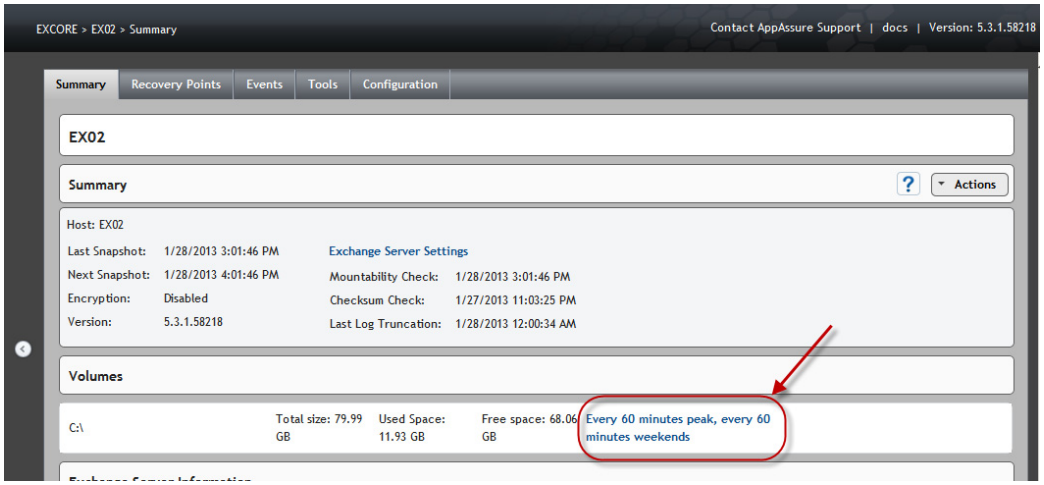

Figure 10. Summary tab - protection schedule

The Edit Protection Settings dialog box displays.

4. In the Edit Protection Settings dialog box, click the **Edit** link next to the volume for which you want to customize the protection schedule.

   The Protection Schedule dialog box displays.

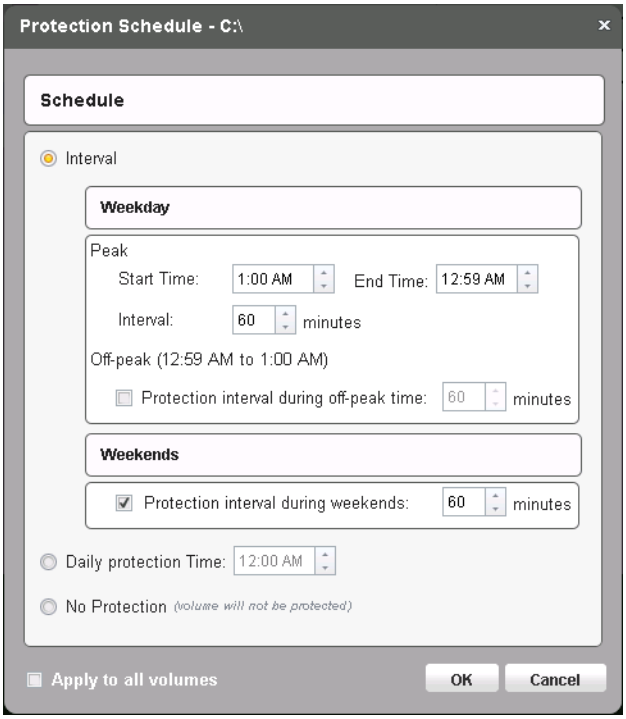

Figure 11. Protection Schedule dialog box

5. In the Protection Schedule dialog box, edit the following schedule options as needed for protecting your data.

   The following table describes the options.

Text Box	Description
Interval	**Weekday**. To protect data on a specific time interval (for example, every 15 minutes), select the **Interval**, and then: □ To customize when to protect data during peak times, you can select a Start Time, End Time, and an Interval from the drop-down menus. □ To protect data during off-peak times, select the **Protection interval during off-peak times** check box, and then select an interval for protection from the drop-down menu. **Weekends**. To protect data during weekends as well, select the **Protection interval during weekends** check box, and then select an interval from the drop-down menu. **NOTE:** If the SQL or Exchange databases and logs are on different volumes, the volumes must belong to one protection group.
Daily	To protect data on a daily basis, select the **Daily** option, and then and in the Protection Time drop-down menu, select a time to start protecting data.
No Protection	To remove protection from this volume, select the **No Protection** option.

   If you want to apply these custom settings to all the volumes on this machine, select **Apply to All Volumes**.

6. When you have made all necessary changes, click **OK**.

Modifying Transfer Settings

In AppAssure 5, 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" on page 33.

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

There are three types of transfers in AppAssure 5:

- **Snapshot.** Backs up the data on your protected machine. Two types of snapshots are possible: a base image of all protected data, and an incremental snapshot for data updated since the last snapshot. This type of transfer creates recovery points, which are stored on the repository associated with the Core.

- **Virtual Machine Export.** Creates a virtual machine (VM) from a recovery point, containing all of the data from the backup of the protected machine, as well the operating system and drivers and associated data to ensure the VM is bootable.
- **Rollback.** Restores backup information to a protected machine.

  NOTE: The entire volume is always rewritten during rollback of Windows systems using EFI system partitions.

Data transfer in AppAssure 5 involves the transmission of a volume of data along a network from AppAssure 5 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. These are some factors that affect data transfer performance:

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

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

To modify transfer settings

1. On the AppAssure 5 Core Console, do one of the following:
   - Click the Machines tab, and then click the hyperlink for the machine you want to modify.
   - Or, in the navigation pane, click the machine you want to modify.
2. Click the Configuration tab, and then click **Transfer Settings**.

   The current transfer settings are displayed.
3. On the Transfer Settings page, click **Change**.

   The Transfer Settings dialog box displays.

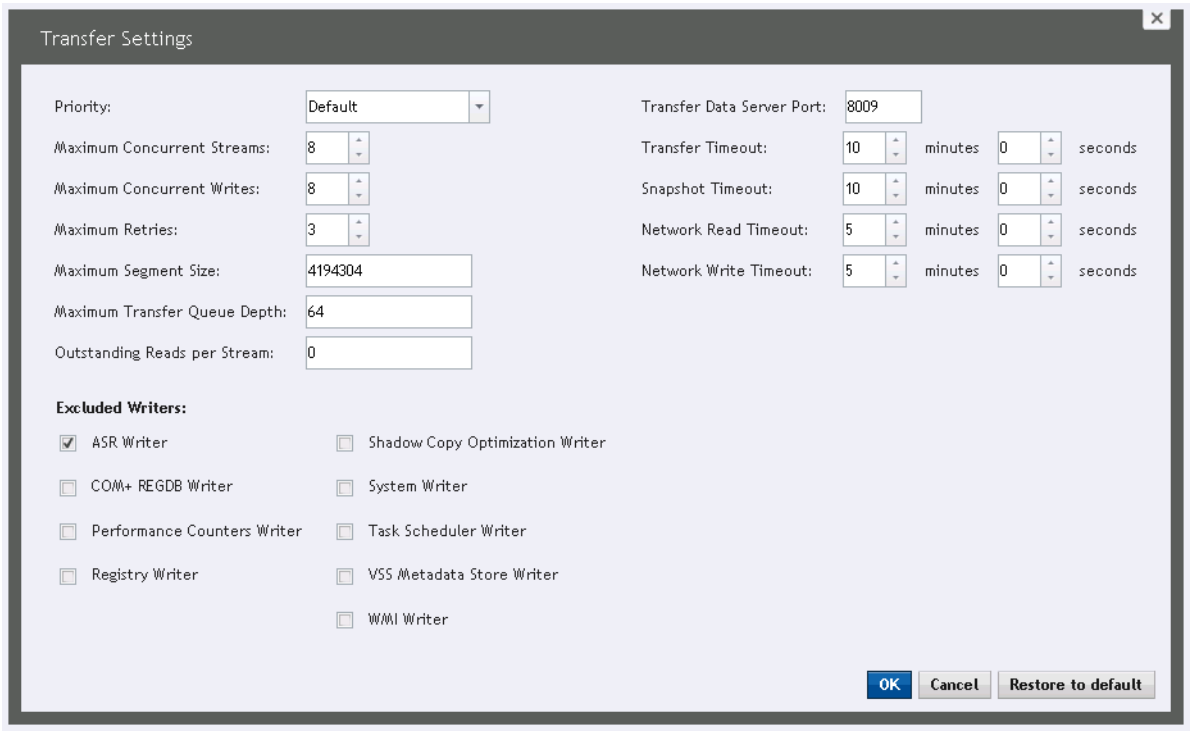

Figure 12. Modify transfer settings

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

Text Box	Description
Priority	Sets the transfer priority between protected machines. Enables you to assign priority by comparison with other protected machines. Select a number from 1 to 10, with 1 being the highest priority. The default setting establishes a priority of 5. **NOTE:** Priority is applied to transfers that are in the queue.
Maximum Concurrent Streams	Sets the maximum number of TCP links that are sent to the Core to be processed in parallel per agent. **NOTE:** Dell recommends setting this value to 8. If you experience dropped packets, try increasing this setting.
Maximum Concurrent Writes	Sets the maximum number of simultaneous disk write actions per agent connection. **NOTE:** Dell recommends setting this to the same value you select for Maximum Concurrent Streams. If you experience packet loss, set slightly lower—for example, if Maximum Current Streams is 8, set this to 7.
Maximum Retries	Sets the maximum number of retries for each protected machine, if some of the operations fail to complete.
Maximum Segment Size	Specifies the largest amount of data, in bytes, that a computer can receive in a single TCP segment. The default setting is 4194304. **CAUTION:** Do not change this setting from the default.
Maximum Transfer Queue Depth	Specifies the amount of commands that can be sent concurrently. You can adjust this to a higher number if your system has a high number of concurrent input/output operations.
Outstanding Reads per Stream	Specifies how many queued read operations will be stored on the back end. This setting helps to control the queuing of agents. **NOTE:** Dell recommends setting this value to 24.

Text Box	Description
Excluded Writers	Select a writer if you want to exclude it. Since the writers that appear in the list are specific to the machine you are configuring, you will not see all writers in your list. For example, some writers you may see include: □ ASR Writer □ BITS Writer □ COM+ REGDB Writer □ Performance Counters Writer □ Registry Writer □ Shadow Copy Optimization Writer □ 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 cannot be performed in that time, the operation is retried.
Network Write Timeout	Specifies the maximum time in seconds to wait for a write connection. If the network write cannot be performed in that time, the operation is retried.

5. Click **OK**.

Restarting a Service

Complete the steps in this procedure to restart a service.

To restart a service

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to restart.
   - Or, in the navigation pane, select the machine you want to restart.
3. Click the Tools tab, and then click **Diagnostics**.
4. Select the **Restart Service** option, and then click the **Restart Service** button.

Viewing Machine Logs

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

To view machine logs

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to view.
   - Or, in the navigation pane, select the machine you want to view.
3. Click the Tools tab, and then click **Diagnostics**.
4. Click the **View Log** link.

Deploying an Agent (Push Install)

AppAssure 5 lets you deploy the AppAssure 5 Agent Installer to individual Windows machines for protection. Complete the steps in the following procedure to push the installer to an agent.

To deploy agents to multiple machines at the same time, see "Deploying to Multiple Machines" on page 122.

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

To deploy an agent

1. Navigate to the AppAssure 5 Core Console and click the Machines tab.
2. In the **Actions** drop-down menu, click **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 agent machine that you want to deploy.
User name	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 should start upon the completion of the deployment and installation of the AppAssure 5 Agent Installer.

4. Click **Verify** to validate the credentials you entered.

   The Deploy Agent dialog box displays a message to indicate that validation is being performed. Click **Abort** if you want to cancel the verification process. After the verification process is complete, a message indicating that verification has been completed displays.

5. Click **Deploy**.

   A message indicating that the deployment has started displays. You can view the progress in the Events tab. Click **Show details** to view more information about the status of the agent deployment.

6. Click **OK**.

Replicating a New Agent

When you add an AppAssure 5 Agent for protection on a source core, AppAssure 5 gives you the option to replicate a new agent to an existing target core. Complete the instructions below to replicate a new agent.

For more information about replication, see "Understanding Replication" on page 48 and "Replicating to a Self-Managed Core" on page 53.

To replicate a new agent

1. Navigate to the AppAssure 5 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 5 Core should use to communicate with the Agent on the machine.
User name	Enter the user name 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 agent machine to be displayed in the AppAssure 5 Core Console.
Repository	Select the repository on the AppAssure 5 Core where the data from this machine should be stored.

Text Box	Description
Encryption Key	Specify whether encryption should be 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 AppAssure 5 Core Console.
Remote Core	Specify a target core to which you want to replicate the agent machine.
Remote Repository	The name of the 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 5 takes a base image.
Schedule	Select one of the following options: □ Protect all volumes with default schedule □ Protect specific volumes with custom schedule **NOTE:** The default schedule is every 15 minutes.For information about custom schedules, see "Creating Custom Schedules for Volumes" on page 96.
Initially pause protection	Select this check box if you want to pause protection, for example, to prevent AppAssure 5 from taking the base image until after peak usage hours.

6. Click **Protect.**

Managing Machines

This section describes a variety of tasks you can perform in managing your machines, such as removing a machine from your AppAssure environment, setting up replication, forcing log truncation, canceling operations, and more.

Removing a Machine

Complete the steps in the following procedure to remove a machine from protection in your AppAssure environment.

To remove a machine

1. Navigate to the AppAssure 5 Core Console and then click the Machines tab.
2. From the Machines tab, perform one of the following:
   - Click the hyperlink for the machine you want to remove.
   - Or, in the navigation pane, select the machine you want to remove.

3. In the **Actions** drop-down menu, click **Remove Machines**, and then select one of the option described in the following table.

Option	Description
Keep Recovery Points	Keeps all currently stored recovery points for this machine.
Remove Recovery Points	Removes all currently stored recovery points for this machine from the repository.

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 setup 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.

For more information about replication, see "Understanding Replication" on page 48 and "Replicating to a Self-Managed Core" on page 53.

To replicate agent data on a machine

1. Navigate to the AppAssure 5 Core Console and then 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 displays.
4. In the Host text box, enter a host name.
5. Under Agents, select the machine that has data 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, in the **Actions** drop-down menu, click **Replication**, and then click **Pause** or **Resume** as needed.

Setting Replication Priority for an Agent

Complete the steps below to set the priority for when an agent replicates.

To set replication priority for an agent

1. From the AppAssure 5 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, an agent with the Highest priority replicates before an agent with the 1 priority.
3. Click **OK.**

Canceling Operations on a Machine

You can cancel currently executing operations for a machine. You can specify to cancel just a current snapshot or to cancel all current operations, which would include exports, replications, and so on.

To cancel operations on a machine

1. From the AppAssure 5 Core Console, click the Machines tab.
2. Select the machine for which you want to cancel operations.
3. In the **Actions** drop-down menu, click **Cancel**, and then select one of the options described in the following table.

Option	Description
All Operations	Cancels all active operations for that machine.
Snapshot	Cancels the snapshot currently in progress.

Viewing Machine Status and Other Details

Complete the step in this procedure to view the status as well as other details for a machine.

To view machine status and other details

- Open the AppAssure 5 Core Console, and do one of the following:
  - Click the Machines tab, and then click the hyperlink for the machine you want to view.
  - Or, in the navigation pane, click the machine you want to view.

  The Summary tab appears.

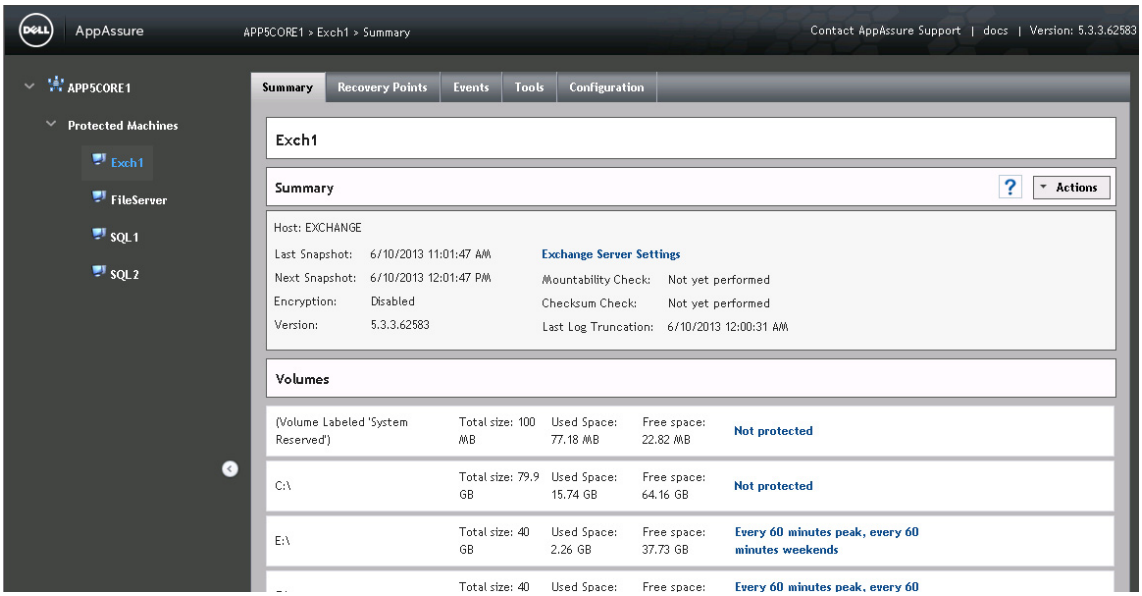

Figure 13. Summary tab - machine status and other information

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

Details information about the volumes contained on this machine also displays and includes:

- Total size

- Used Space
- Free space

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

- Name
- Install Path
- Version
- Version Number
- Database Name
- Online status

If Exchange Server is installed on the machine, detailed information about the server and mail stores also displays 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 5 Agent software simultaneously to multiple Windows machines.

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

1. Deploy AppAssure 5 to multiple machines. See "Deploying to Multiple Machines" on page 122.
2. Monitor the activity of the batch deployment. See "Monitoring the Deployment of Multiple Machines" on page 128.
3. Protect multiple machines. See "Protecting Multiple Machines" on page 129.

   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" on page 131.

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 5. From within the Core Console, you can specifically bulk deploy to:

- Machines on an Active Directory domain
- Machines on a VMware vCenter/ESX(i) virtual host
- 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.

You can use the Bulk Deploy feature to deploy the Agent software to as many as 50 agent machines. The machines to which you are deploying must have access to the Internet to download and install bits as AppAssure 5 uses the Web version of the AppAssure 5 Agent Installer to deploy the installation components. If access to the Internet is not available, you will need to manually download the AppAssure 5 Agent Installer from the License Portal and deploy the installer to the machines.

For more information, see "Managing AppAssure 5 Licenses", on the AppAssure 5 Technical Documentation page at: http://docs.appassure.com/display/AA50D/AppAssure+5+Technical+Documentation.

Deploying to Machines on an Active Directory Domain

Before starting this procedure, you must have the domain information and logon credentials for the Active Directory server.

To deploy to multiple machines on an Active Directory domain

1. On the AppAssure 5 Core Console, click the Tools tab, and then click **Bulk Deploy**.
2. On the Deploy Agent to Machines window, click **Active Directory**.
3. In the Connect to Active Directory dialog box, enter the domain information and logon credentials as described in the following table.

Text Box	Description
Domain	The host name or IP address of the Active Directory domain.
User name	The user name used to connect to the domain; for example, Administrator.
Password	The secure password used to connect to the domain.

4. Click **Connect**.

5. On the Add Machines from Active Directory dialog box, select the machines to which you want to deploy the AppAssure 5 Agent, and then click **Add**.

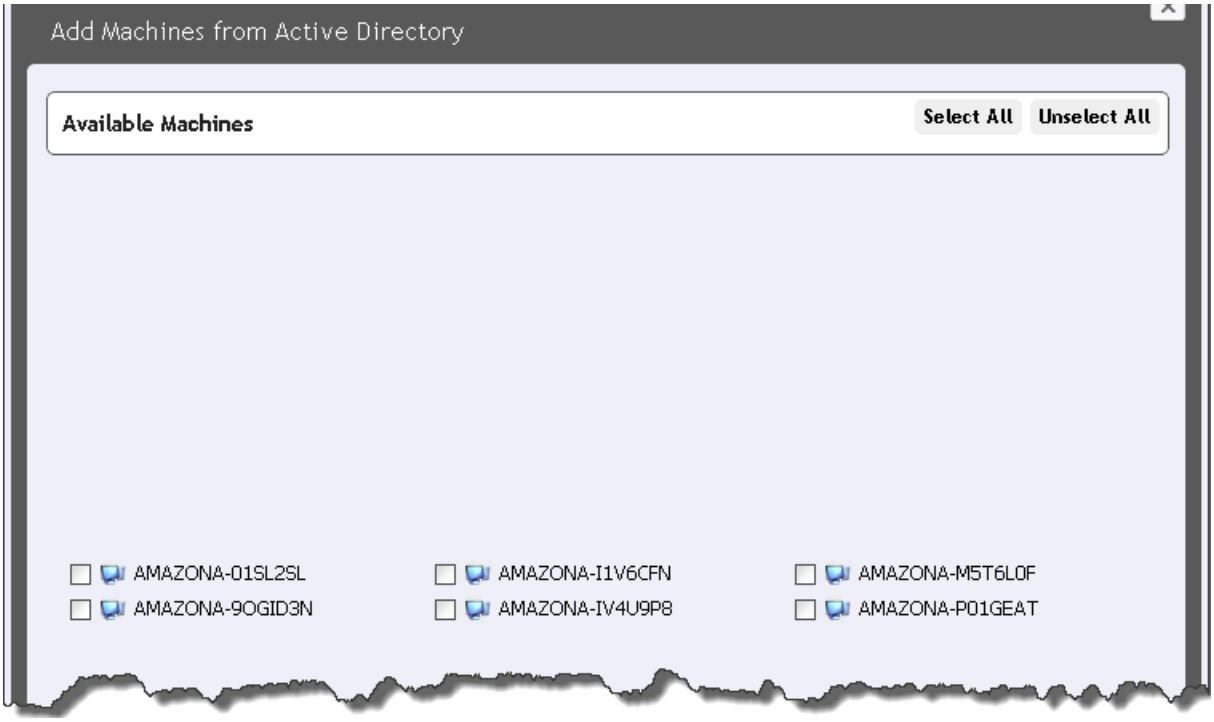

Figure 14. Add Machines from Active Directory dialog box

The machines you added appear on the Deploy Agent to Machines window.

6. To enter the password for the machine, select a repository, add an encryption key, or edit other settings for a machine, click the **Edit** link for that machine, and then do the following:

   a. In the Edit Settings dialog box, specify the settings as described in the following table.

Text Box	Description
Host name	Automatically provided from step 3.
Display name	Automatically assigned based on the host name provided in step 3.
Port	The port number on which the AppAssure 5 Core communicates with the agent on the machine.
User name	Automatically provided from step 3.
Password	Enter the password for the machine.

Text Box	Description
Automatic reboot after install	Specify whether you want to automatically reboot the machine after deployment. **NOTE:** This option is mandatory and selected by default if you want to automatically protect the machine after deployment by selecting **Protect machine after install**.
Protect machine after install	Specify whether you want to automatically protect the machine after deployment. (This allows you to skip "Protecting Multiple Machines" on page 129.)
Repository	Use the drop-down list to select the repository on the AppAssure 5 Core where the data from the machines should be stored. The repository you select is used for all machines being protected. **NOTE:** This option is only available when you select **Protect machine after install**.
Encryption Key	(Optionally) Use the drop-down list to specify whether encryption should be applied to the data on the machine that should be stored in the repository. The encryption key is assigned to all machines that are being protected. **NOTE:** This option is only available when you select **Protect machine after install**.

b. Click **Save**.

7. To verify that AppAssure 5 can connect to each machine successfully, select each machine in the Deploy Agent to Machines window, and then click **Verify**.

   The Deploy Agent on Machines window shows an icon next to each machine that reflects its readiness for deployment, as follows:

   - Green icon - AppAssure 5 is able to connect to the machine and it is ready to be deployed.
   - Yellow icon - AppAssure 5 is able to connect to the machine; however, the AppAssure 5 Agent on the machine is already paired with an AppAssure 5 Core.
   - Red icon - AppAssure 5 cannot connect to the machine. This may be because the logon credentials are incorrect, the machine is shut down, the firewall is blocking traffic, or another problem. To correct, click **Edit Settings** on the toolbar or the **Edit** link next to the machine.

8. After machines are successfully verified, select each machine to which you want to deploy the AppAssure 5 Agent, and then click **Deploy**.

   If you chose the **Protect machine after install** option, after deployment is successful, the machines automatically reboot and protection is enabled.

Deploying to Machines on a VMware vCenter/ESX(i) Virtual Host

Before starting this procedure, you must have the host location information and logon credentials for the VMware vCenter/ESX(i) virtual host.

Note

All virtual machines must have VM Tools installed; otherwise, AppAssure 5 cannot detect the host name of the virtual machine to which to deploy. In lieu of the host name, AppAssure 5 uses the virtual machine name, which may cause issues if the host name is different from the virtual machine name.

To deploy to multiple machines on a vCenter/ESX(i) virtual host

1. Navigate to the AppAssure 5 Core Console, click the Tools tab, and then click **Bulk Deploy**.
2. In the Deploy Agent to Machines window, click **vCenter/ESX(i)**.
3. In the Connect to vCenter Server/ESX(i) dialog box, enter the host information and logon credentials as follows and click **Connect**.

Text Box	Description
Host	The name or IP address of the VMware vCenter Server/ESX(i) virtual host.
User name	The user name used to connect to the virtual host; for example, Administrator.
Password	The secure password used to connect to this virtual host.

4. In the Add Machines from VMware vCenter Server/ESX(i) dialog box, select the machines to which you want to deploy and then click **Add**. (You can expand and drill through the machines by clicking the arrow icon next to the machine name.)

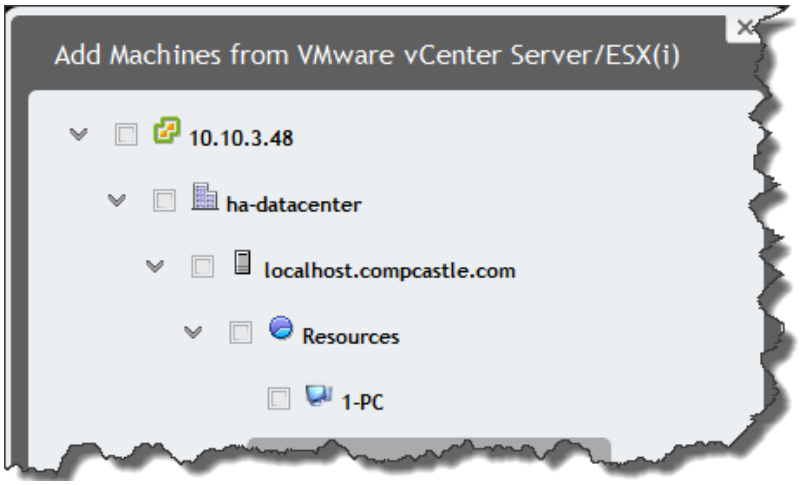

Figure 15. Add Machines from VMware vCenter Server dialog box

5. In the Deploy Agent on Machines window, you should see the machines that you added. If you want to select a repository, encryption key, or other settings for a machine, such as automatically rebooting the machine once the software is installed, select the machine and click **Edit**.

   For details on each setting, see "Deploying to Machines on an Active Directory Domain" on page 122.

6. Verify that AppAssure 5 can connect to each machine successfully. To do this, select each machine in the Deploy Agent on Machines window, and click **Verify**.

   The Deploy Agent on Machines window shows an icon next to each machine that reflects its readiness for deployment, as follows:

   - Green icon - AppAssure 5 is able to connect to the machine and it is ready to be deployed.
   - Yellow icon - AppAssure 5 is able to connect to the machine; however, the agent is already paired with a core machine.
   - Red icon - AppAssure 5 cannot connect to the machine. This may be because the logon credentials are incorrect, the machine is shut down, the firewall is blocking traffic, or another problem. To correct, click **Settings** on the toolbar or the **Edit** link next to the machine.

7. After machines are verified successfully, select each machine and click **Deploy**.

   If you chose the **Protect machine after install** option, after deployment is successful, the machines are rebooted automatically and protection is enabled.

Deploying to Machines on Any Other Host

To deploy to multiple machines on any other host

1. Navigate to the AppAssure 5 Core Console, click the Tools tab, and then click **Bulk Deploy**.

2. On the Deploy Agent on Machines window, do one of the following:

   - Click **New** to enter a new machine host, logon credentials, repository, encryption key, and other information. For details on each setting, see "Deploying to Machines on an Active Directory Domain" on page 122.

     After you enter this information, click **OK** to add it to the Deploy Agent on Machines list, or click **OK & New** to add another machine.

     NOTE: If you want to automatically protect the machine after deployment, check the **Protect Machine after Install** box. If you check the box, the machine will be rebooted automatically prior to enabling protection.

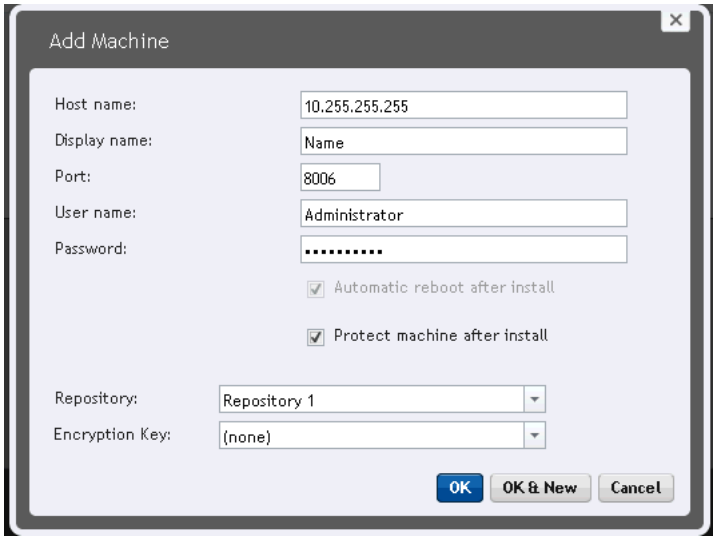

Figure 16. Add Machine dialog box

   - To specify multiple machines in a list, click **Manually**, enter the machine details in the Add Machines Manually dialog box, and click **Add**. For each machine, you will need to enter the IP address or name for the machine, the user name, the password separated by a double-colon delimiter, and port as shown in the following format:

     ```
     hostname::username::password::port
     ```

     For example:

     ```
     10.255.255.255::administrator::&11@yYz90z::8006

     abc-host-00-1::administrator::99!zU$o83r::168
     ```

3. In the Deploy Agent on Machines window, you should see the machines that you added. If you want to select a repository, encryption key, or other settings for a machine, select the machine and click **Edit**.

   For details on each setting, see "Deploying to Machines on an Active Directory Domain" on page 122.

4. Verify that AppAssure 5 can connect to each machine successfully. To do this, select each machine in the Deploy Agent on Machines window, and click **Verify**.

   The Deploy Agent on Machines window shows an icon next to each machine that reflects its readiness for deployment, as follows:

   - Green icon - AppAssure 5 is able to connect to the machine and it is ready to be deployed.
   - Yellow icon - AppAssure 5 is able to connect to the machine; however, the AppAssure 5 Agent on the machine is already paired with an AppAssure 5 Core.
   - Red icon - AppAssure 5 cannot connect to the machine. This may be because the logon credentials are incorrect, the machine is shut down, the firewall is blocking traffic, or another problem. To correct, click **Settings** on the toolbar or the **Edit** link next to the machine.

5. After machines are verified successfully, check the box next to each machine and click **Deploy**.

   If you chose the **Protect machine after install** option, after deployment is successful the machines are rebooted automatically and protection is enabled.

Monitoring the Deployment of Multiple Machines

You can view the progress of the deployment of AppAssure 5 Agent software to the machines. Complete the steps in this procedure to view the deployment.

To monitor the deployment of multiple machines

1. Navigate to the AppAssure 5 Core Console, click the Events tab, select the deployment job in the list, and click the Details icon in the Details column.

   The Monitor Active Task window displays to present the details of the deployment. The details include overall progress information, including:

   - Start Time
   - End Time
   - Status
   - Elapsed Time
   - Time Remaining
   - Progress
   - Phase
   - Failure Reason, if applicable

2. Click **Open in New window** to launch a new window to view the progress of the deployment. Or, click **Close** to close the window; the deployment task will process in the background.

Protecting Multiple Machines

After bulk deploying the AppAssure 5 Agent software to your Windows machines, you will need to protect the machines to protect the data. If you selected **Protect Machine After Install** when you deployed the Agent, you can skip this step.

Note

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

To protect multiple machines

1. Navigate to the AppAssure 5 Core Console, click the Tools tab, and then click **Bulk Protect**.

   The Protect Machines window displays.

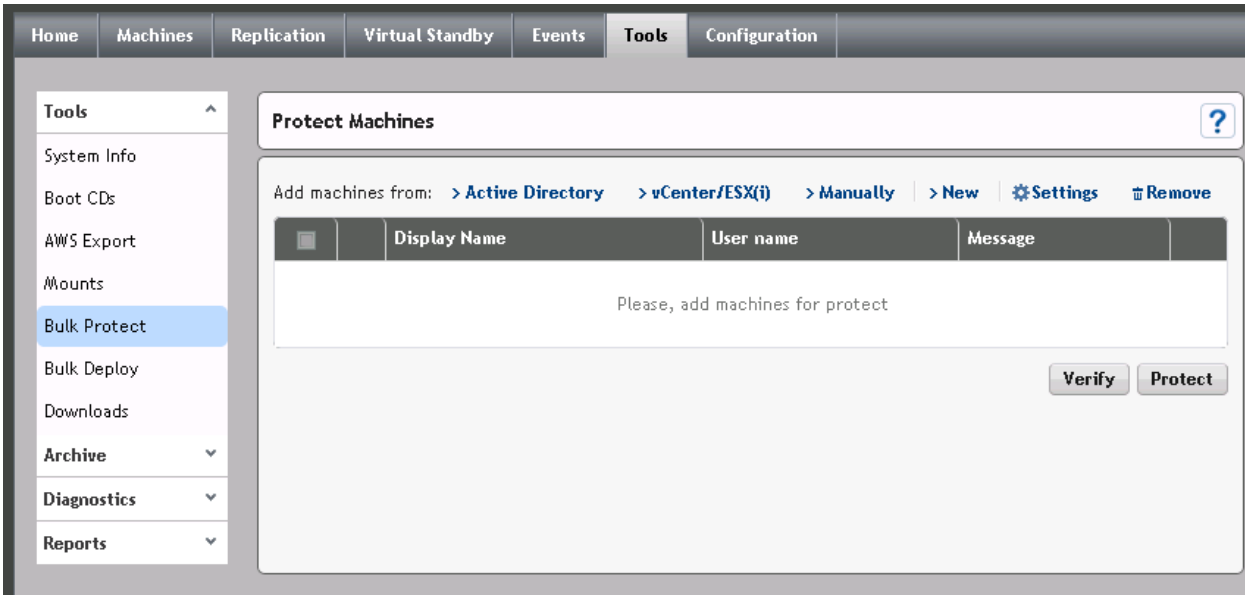

Figure 17. Protect Machines window

2. Add the machines you want to protect by clicking one of the following options. For details on how to add each option for each type of machine, see the related option in the section, "Deploying to Multiple Machines" on page 122.
   - Click **Active Directory** to specify machines on an Active Directory domain.
   - Click **vCenter/ESX(i)** to specify virtual machines on a vCenter/ESX(i) virtual host.
   - Click **New** to specify other types of machines one by one.
   - Click **Manually** to specify multiple machines in a list by entering host name, credentials, and port information in the following format:

     ```
     hostname::username::password::port
     ```

3. Once you have added the machines for protection, in the Protect Machines window, you should see the newly added machines. If you want to select a repository, encryption key, or other advanced settings for a machine, select the machine and click **Edit**.
4. Specify the settings as follows and click **OK**.

Text Box	Description
User name	The user name used to connect to this machine; for example, Administrator.
Password	The secure password used to connect to this machine.
Port	The port number on which the AppAssure 5 Core communicates with the Agent on the machine.
Repository	Select the repository on the AppAssure 5 Core where the data from the machines should be stored. The repository you select is used for all machines being protected.
Encryption Key	Specify whether encryption should be applied to the Agent on the machines that should be stored in the repository. The encryption key is assigned to all machines that are being protected.
Protection Schedule	The schedule for which the protection of the machine occurs. The default schedule is 60 minutes during peak operation and 60 minutes on weekends. □ To edit the schedule to suit the needs of your enterprise, click **Edit**. For more information, see "Modifying Protection Schedules" on page 108.
Initially Pause Protection	Optionally, you can choose to pause protection when first run; that is, the Core will not take snapshots of the machines until you manually resume protection.

5. Verify that AppAssure 5 can connect to each machine successfully. To do this, select each machine in the Protect Machines window, and click **Verify**.

   The Protect Machines window shows an icon next to each machine that reflects its readiness for deployment, as follows:

   - Green icon - AppAssure 5 is able to connect to the machine and it is ready to be protected.
   - Yellow icon - AppAssure 5 is able to connect to the machine; however, the AppAssure 5 Agent on the machine is already paired with an AppAssure 5 Core.
   - Red icon - AppAssure 5 cannot connect to the machine. This may be because the logon credentials are incorrect, the machine is shut down, the firewall is blocking traffic, or another problem. To correct, click **Settings** on the toolbar or the **Edit** link next to the machine.
6. After the machines are verified successfully, select each machine, and click **Protect**.

Monitoring the Protection of Multiple Machines

You can monitor the progress as AppAssure 5 applies the protection polices and schedules to the machines.

To monitor the protection of multiple machines

1. Click the Machines tab to view the status and progress of the protection.

   The Protected Machines page displays.

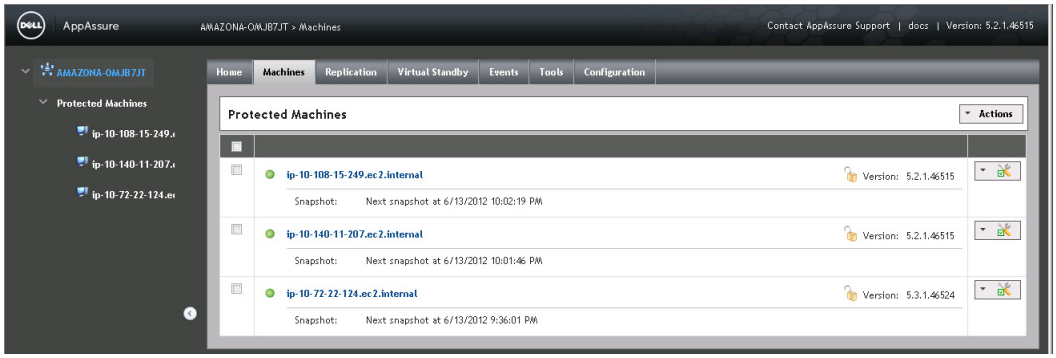

Figure 18. Machines tab - protected machines status

2. Click the Events tab to view related tasks and alerts.

   The Events tab displays, showing Task and Alert events. As volumes are transferred, the status, start times, and end times display in the Tasks pane. As each protected machine is added, an alert is logged on the Events tab, which lists whether the operation was successful or if errors were logged.

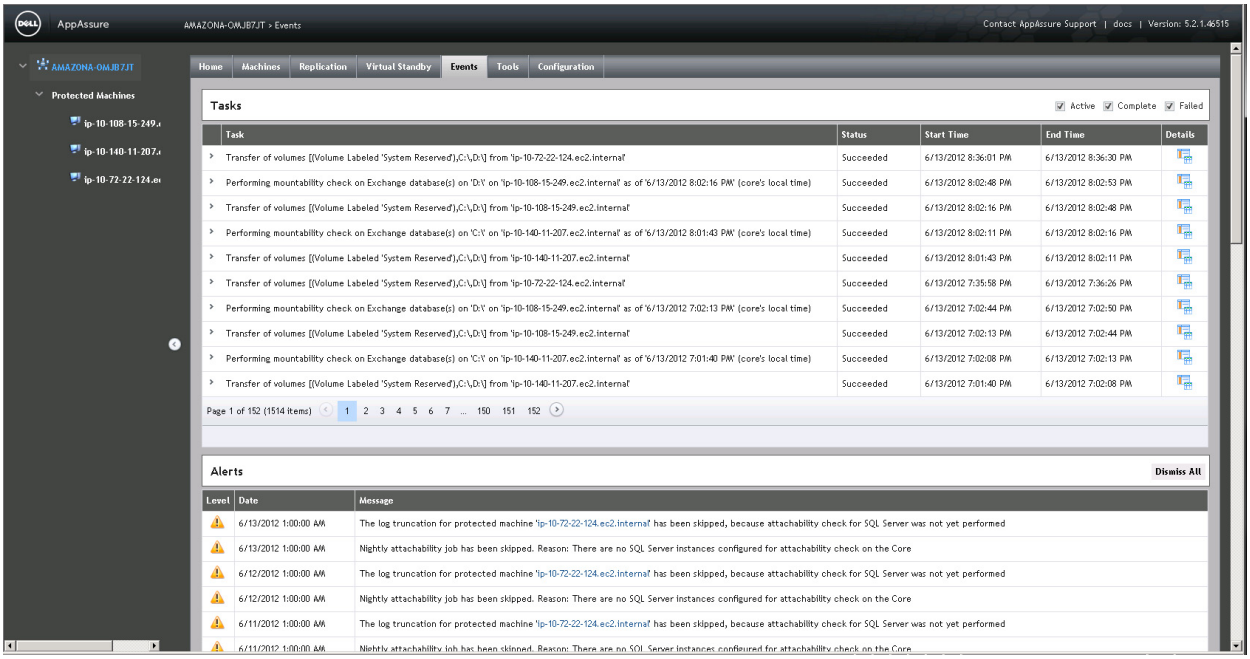

Figure 19. Events tab - Tasks and Alerts

3. In the upper right area of the Events tab, you can control how active and completed events are displayed by doing the following:
   - To show only active events, ensure only **Active** is selected.
   - To show only completed events, ensure that only **Completed** is selected.
   - To show both active and completed events, ensure that both **Active** and **Completed** are selected.

To view task details

1. Select the protection task in the list, and, in the Details column, click the Details icon to view more specific information about the task.

   The Monitor Active Task window displays to present the details of the protection task. The details include overall progress information, including:
   - Start Time
   - End Time
   - Status
   - Elapsed Time
   - Time Remaining

- Progress
- Phase
- Failure Reason, if applicable

2. Click **Open in New window** to launch a new window to view the progress of the protection. Or, click **Close** to close the window; the task will process in the background.

To view alert information

- In the Alerts pane, you will see an alert as each protected machine is added. This alert lists whether the operation was successful, any errors that occurred, the level of the alert, the transactional date, and the related message.

  If you want to remove all alerts from the page, you can click **Dismiss All.**

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 5, you can force a snapshot, temporarily pause snapshots, and view lists of current recovery points in the repository as well as delete them if needed. Recovery points are used to restore protected machines or to mount to a local file system.

The snapshots that are captured by AppAssure 5 are done so at the block level and are application aware. This means that all open transactions and rolling transaction logs are completed and caches are flushed to disk before creating the snapshot.

AppAssure 5 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

Complete the steps in the following procedure to view recovery points.

To view recovery points

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

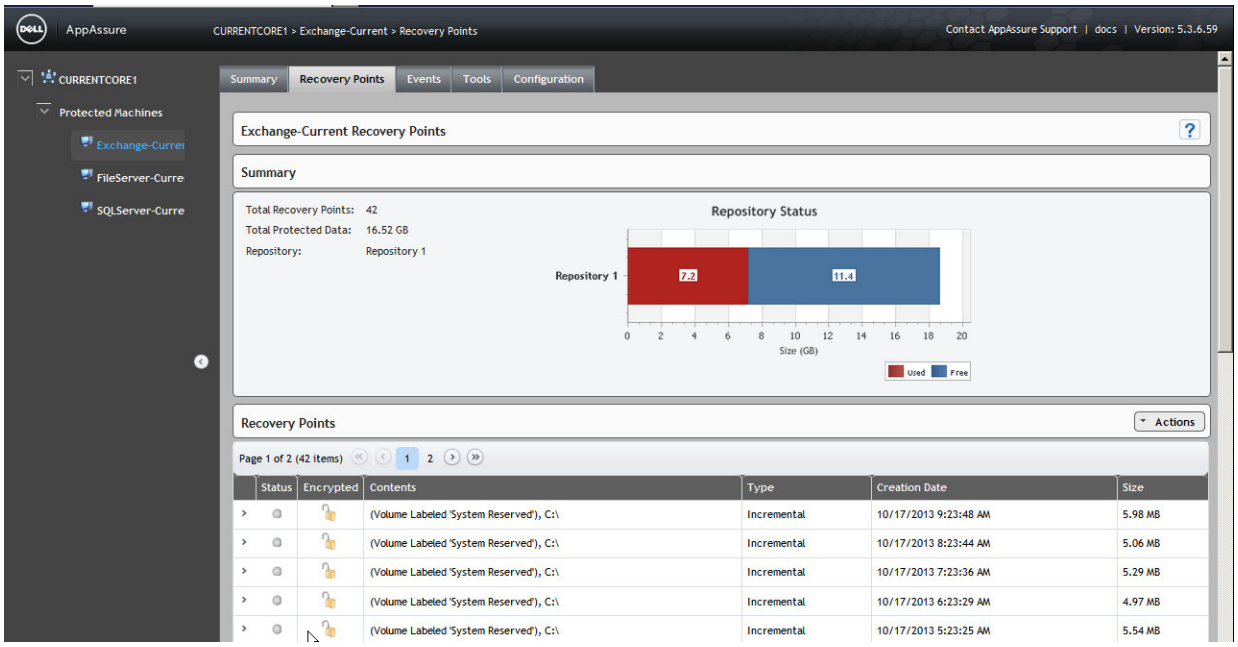

Figure 20. Recovery Points tab

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

Complete the steps in the following procedure to view details about a specific recovery point.

To view a specific recovery point

1. In the left navigation area of the AppAssure Core Console, select the machine for which you want to view recovery points, and then click the Recovery Points tab.
2. Click the right angle bracket **>** symbol next to a recovery point in the list to expand the view.

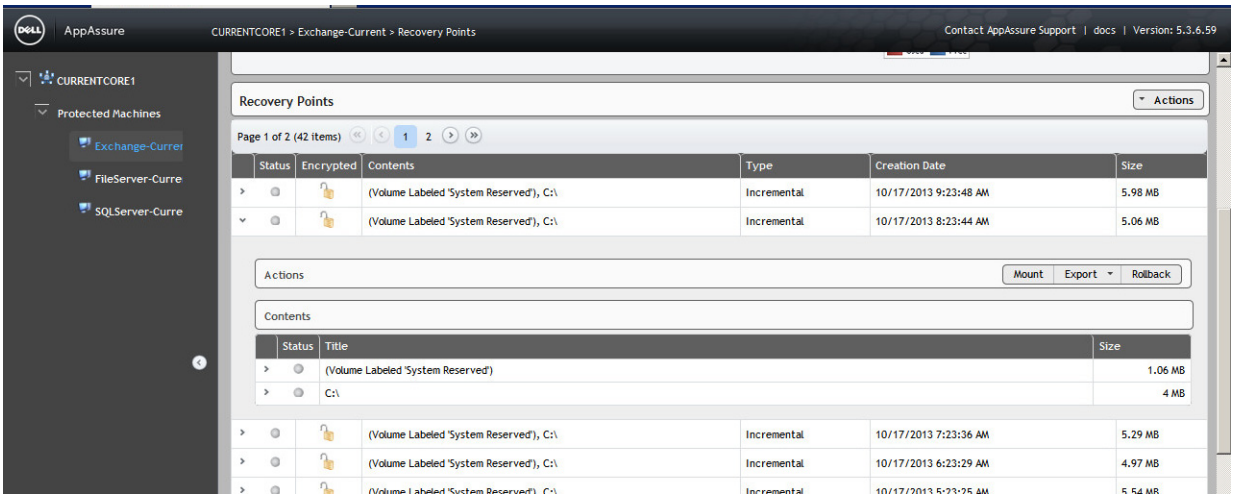

Figure 21. Recovery Points tab – selected recovery point details

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, as described in the following table.

Info	Description
Actions	The Actions menu includes the following operations you can perform on the selected recovery point: **Mount**. Select this option to mount the selected recovery point. For more information, see "Mounting a Recovery Point for a Windows Machine" on page 136. **Export**. From the Export option, you can export the selected recovery point to ESXi, VMware workstation, or HyperV. For more information, see "Exporting Backup Information for your Windows Machine to a Virtual Machine" on page 144. **Rollback**. Select this option to perform a restore from the selected recovery point to a volume you specify. For more information, see "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165.
Contents	The Contents area includes a row for each volume in the expanded recovery point, listing the following information for each volume: **Status**. Indicates current status of the recovery point. **Title**. Lists the specific volume in the recovery point. **Size**. Displays the amount of space that the recovery point consumes in the repository.

3. Click the right angle bracket **>** symbol 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.
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 5, you can mount a recovery point for a Windows machine to access stored data through a local file system.

When mounting recovery points from data restored from Windows machines that has data deduplication enabled, you will need to make sure that deduplication is also enabled on the Core server.

To mount a recovery point for a Windows machine

1. In the AppAssure 5 Core Console, navigate to the Mount Recovery Point dialog box by doing one of the following:
   - Click the Machines tab.
     i. In the **Actions** drop-down menu for the machine or cluster with the recovery that you want to mount, click **Mount**.
     ii. In the Mount Recovery Point dialog box, click to select a recovery point in the list, and then click **Next**.

     The Mount Recovery Points dialog box appears.
   - In the left navigation area of the AppAssure 5 Core Console, select the machine that you want to mount to a local file system.

     The Summary tab for the selected machine displays.
     i. Click the Recovery Points tab.
     ii. In the list of recovery points, click the right angle bracket **>** symbol to expand the recovery point that you want to mount.
     iii. In the expanded details for that recovery point, click **Mount.**

     The Mount Recovery Points dialog box appears.

2. In the Mount Recovery Point dialog box, edit the text boxes for mounting a recovery point as described in the following table.

Option	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: □ 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.

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

   The Active Task dialog box appears.

4. In the Active Task dialog box, click **Open Monitor Window** to monitor the task for mounting the selected recovery point, or click **Close** to dismiss this dialog box.

   NOTE: For more information about monitoring AppAssure 5 events, see "Viewing Events and Alerts" on page 186.

Dismounting Select Recovery Points

Complete the steps in this procedure to dismount select recovery points that are mounted locally on the Core.

To dismount select recovery points

1. Navigate to the AppAssure 5 Core, and then click the Tools tab.
2. From the Tools option, click **System Info**.
3. Locate and select the mounted display for the recovery point you want to dismount, and then click **Dismount**.

Dismounting All Recovery Points

Complete the steps in this procedure to dismount all recovery points that are mounted locally on the Core.

To dismount all recovery points

1. Navigate to the AppAssure 5 Core, and then click the Tools tab.

2. From the Tools option, click **System Info**.
3. In the Local Mounts section, click **Dismount All**.

Mounting a Recovery Point Volume on a Linux Machine

In AppAssure you can remotely mount a restored volume to the local disk from a recovery point for a Linux machine.


Note

When performing this procedure, do not attempt to mount recovery points to the /tmp folder, which contains the aavdisk files.

To mount a recovery point 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 AppAssure Core server.
6. Enter the logon credentials for the Core server, that is, the user name and password.

   A list displays showing the machines protected by this AppAssure 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. Enter the following command to list the currently mounted recovery points for a specified machine:

   ```
   lr <line_number_of_machine>
   ```

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

   A list displays 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. 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>

   NOTE: You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, you would use the agent/machine line number (from the lm output), followed by the recovery point line number and volume letter, followed by the path, such as, m <machine_line_number> <recovery_point_line_number> <volume_letter> <path>. For example, if the `lm` output lists three agent machines, and you enter the `lr` command for number 2 and you to mount the 23 recovery point volume b to /tmp/mount_dir the command would be:
   m 2 23 b /tmp/mount_dir

9. To verify the mount was successful, enter the following command, which should list the attached remote volume:

   `l`

Caution

Note that you should not unmount a protected Linux volume by hand. In the event you need to do this, 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 would need to be in a form similar to this example:
/dev/sda1.

Removing Recovery Points

You can easily remove recovery points for a particular machine from the repository. When you delete recovery points in AppAssure 5, you can specify one of the following options.

- **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 AppAssure 5 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.

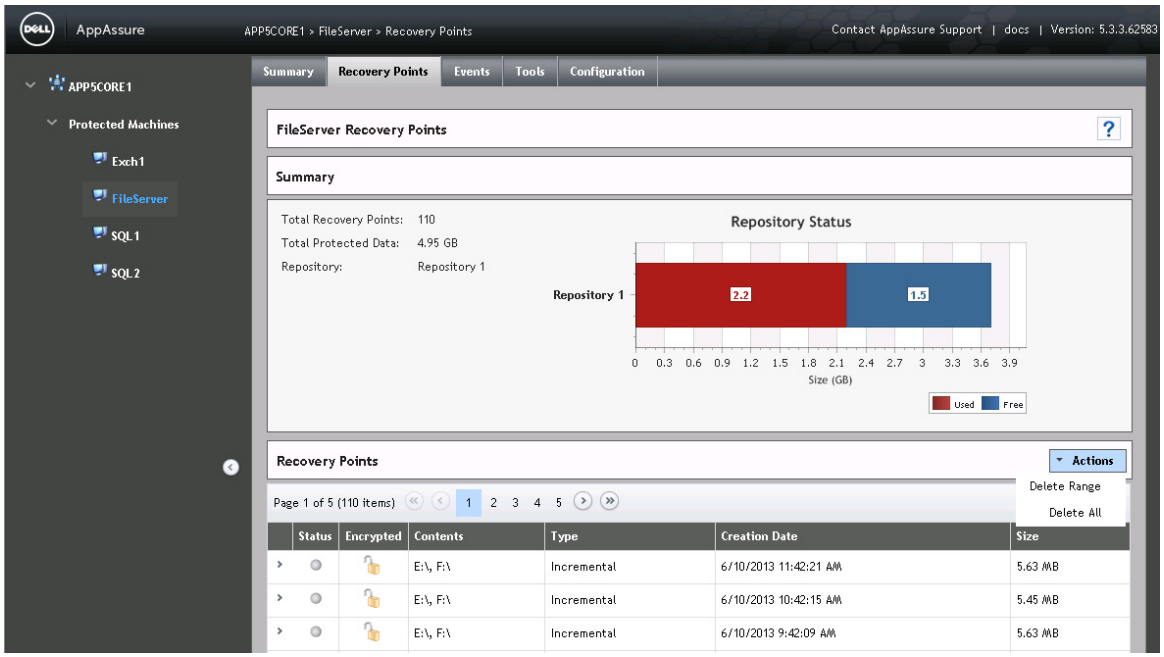

Figure 22. Recovery Points tab - machine status and other information

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 displays.

     - In the Delete Range dialog box, specify the range of recovery points you want to delete 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; however, without the base image, the resulting recovery points are incomplete and are unlikely to contain the data necessary to complete a recovery. These recovery points are considered to be part of the orphaned recovery point chain. If this situation occurs, the best solution is to delete the chain and create a new base image.

For more information about forcing a base image, see "Forcing a Snapshot" on page 141.

Note

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

To delete an orphaned recovery point chain

1. On the AppAssure 5 Core Console, click 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 windows appears.
5. In the Delete Recovery Points window, click **Yes**.

Caution

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

The orphaned recovery point chain is deleted.

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, which is referred to as a base image.

To force a snapshot

1. In the AppAssure 5 Core Console, click the Machines tab, and then, in the list of protected machines, select the machine or cluster with the recovery point for which you want to force a snapshot.

2. Click the **Actions** drop-down menu for that machine, click **Force Snapshot**, and then select one of the options described in the following table.

Option	Description
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 notification appears in the Transfer Status dialog box that the snapshot has been queued, click **OK**.

   A progress bar displays next to the machine in the Machines tab to show the progress of the snapshot.

Restoring Data

Using the Live Recovery instant recovery technology in AppAssure 5, 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. The topics in this section describe how you can export protected data from a Windows machine to a virtual machine or roll back a Windows or Linux 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. For details, see "Replicating Agent Data on a Machine" on page 118.

Windows 8, 8.1 and Windows Server 2012, 2012 R2 operating systems that are booted from FAT32 EFI partitions are not available for protection or recovery, nor are Resilient File System (ReFS) volumes. For details, see the AppAssure 5 Deployment Guide.

When recovering data on Windows machines, if the volume that you are restoring has Windows data deduplication enabled, you will need to make sure that deduplication is also enabled on the Core server.

About Exporting Protected Data from Windows Machines to Virtual Machines

AppAssure 5 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.

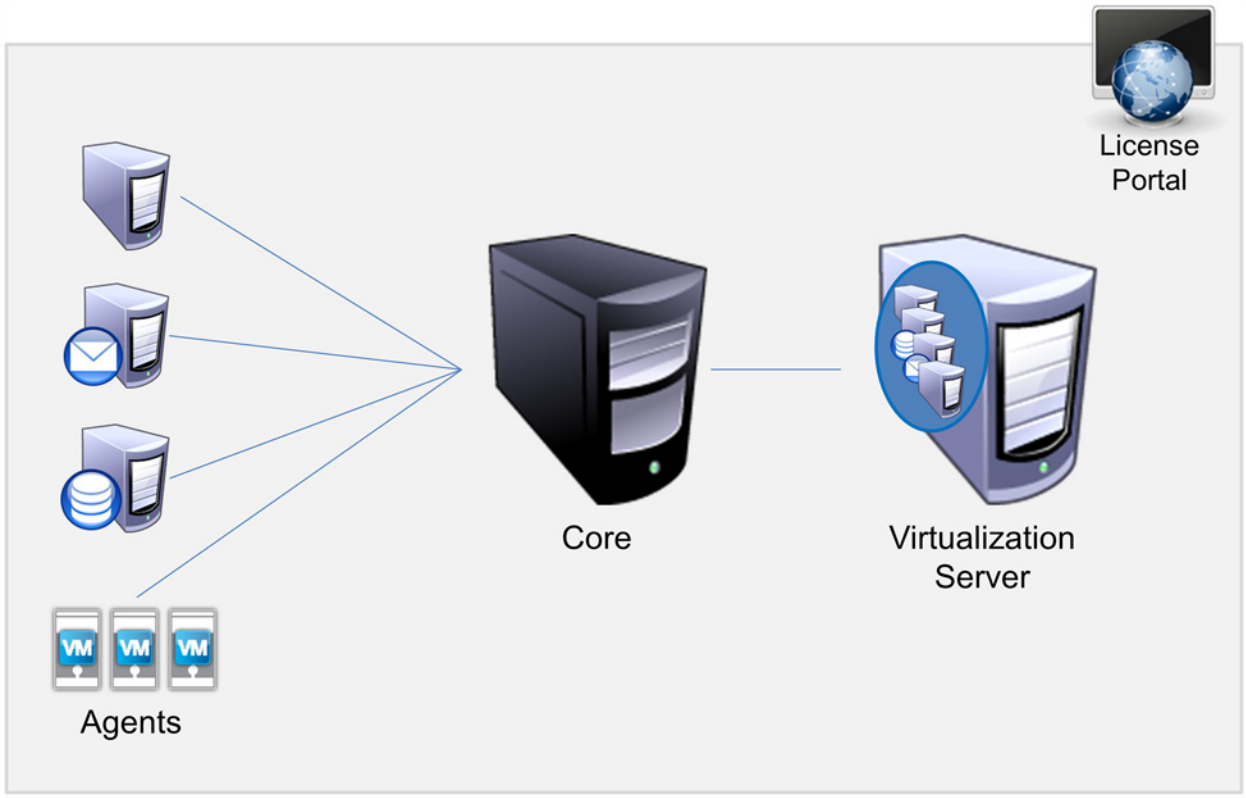

Figure 23. Virtual standby deployment

You create a virtual standby by continuously exporting protected data from your Windows machine to a virtual machine (VMware, ESXi, and Hyper-V). 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.

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

AppAssure 4.x and 5.x both support taking snapshots of all dynamic and basic volumes. AppAssure 4.x and 5.x also support exporting simple dynamic volumes that are on a single physical disk. As their name implies, simple dynamic volumes are not striped, mirrored or spanned volumes.

Non-simple dynamic volumes have arbitrary disk geometries that cannot be fully interpreted and therefore AppAssure cannot export them. Neither Replay 4.x nor AppAssure 5.x has the ability to export complex or non-simple dynamic volumes.

AppAssure does not support exporting non-simple or complex dynamic volumes. Notification appears in the user interface to alert you that exports are limited and restricted to simple dynamic volumes. If you attempt to export anything other than a single dynamic volume, the export job will fail.

Exporting Backup Information for your Windows Machine to a Virtual Machine

In AppAssure 5 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

In AppAssure 5, Version 5.3.6.125 a limitation exists when attempting to perform a VM export on machines that have Windows 8.1 or Windows Server 2012 R2 installed. For more information about the limitation, see the AppAssure 5 Release Notes for 5.3.6.125 and the AppAssure 5 Knowledge Base.

To export Windows backup information to a virtual machine

- In the AppAssure 5 Core Console, click the Machines tab, and then do the following:

  a. In the list of protected machines, select the machine or cluster with the recovery point for which you want to export.

  b. 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:

     - ESXi Export
     - VMware Workstation Export
     - Hyper-V Export

  The Select Export Type dialog box displays.

Exporting Windows Data using ESXi Export

In AppAssure 5, you can choose to export data using ESXi Export by performing a one-time or continuous export. Complete the steps in the following procedures to export using ESXi Export for the appropriate type of export.

Performing a One-Time ESXi Export

You can choose to perform a one-time export for ESXi. Complete the steps in this procedure to perform a one-time export.

To perform a one-time ESXi export

1. In the AppAssure 5 Core Console, click the Machines tab.
2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **ESX(i) Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **One-time export** and click **Next**.

   The ESXi Export - Select Recovery Point dialog box displays.
5. Select a recovery point to export and then click **Next.**

   The Virtual Standby Recovery Point to VMware vCenter Server/ESXi dialog box displays.

Defining Virtual Machine Information for Performing an ESXi Export

Complete the steps in this procedure to define the information for the virtual machine.

To define virtual machine information for performing an ESXi export

1. From the Virtual Standby Recovery Point to VMware vCenter Server/ESXi dialog box, enter the parameters for accessing the virtual machine as described in the following table.

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.

2. Click **Connect**.

Performing a Continuous (Virtual Standby) ESXi Export

You can choose to perform a one-time or continuous export. Complete the steps in this procedure to perform a continuous export.

To perform a continuous (virtual standby) ESXi export

1. In the AppAssure 5 Core Console, click the Machines tab.

2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **ESX(i) Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **Continuous (Virtual Standby)**.
5. Click **Next**.

   The Virtual Standby Recovery Point to VMware vCenter Server/ESXi dialog box displays.
6. Enter the parameters for accessing the virtual machine as described in the following table.

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.

7. Click **Connect**.
8. In the Options tab, enter the information for the virtual machine as described in the following table.

Option	Description
Virtual Machine Name	Enter a name for the virtual machine.
Memory	Specify the memory usage. You can choose from the following options: □ Use the same amount of RAM as source machine □ Use a specific amount of RAM, and then specify the amount in MB
ESXi Datacenter	Enter the name for the ESXi data center.
ESXi Host	Enter the credentials for the ESXi host.
Data Store	Enter the details for the data store.
Resource Pool	Enter a name for the resource pool.

9. Click **Start Export**.

Exporting Windows Data using VMware Workstation Export

In AppAssure 5, you can choose to export data using VMware Workstation Export by performing a one-time 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

You can choose to perform a one-time export for VMware Workstation Export. Complete the steps in this procedure to perform a one-time export.

To perform a one-time VMware Workstation export

1. In the AppAssure 5 Core Console, click the Machines tab.
2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **VMware Workstation Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **One-time export**.
5. Click **Next**.

   The VM Export - Select Recovery Point dialog box displays.
6. Select a recovery point to export and then click **Next.**

   The Virtual Standby Recovery Point to VMware Workstation/Server dialog box displays.

Defining One-Time Settings for Performing a VMware Workstation Export

Complete the steps in this procedure to define the settings for performing a one-time VMware Workstation export.

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

1. From the Virtual Standby Recovery Point to VMware Workstation/Server dialog box, enter the parameters for accessing the virtual machine as described in the following table.

Option	Description
Target Path	Specify the path of the local folder or network share on which to create the virtual machine. **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.
User name	Enter the logon credentials for the virtual machine. □ If you specified a network share path, you need to 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	Enter the logon credentials for the virtual machine. □ If you specified a network share path, you need to 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.

2. In the Export Volumes pane, select the volumes to export; for example, C:\ and D:\.
3. In the Options pane, enter the information for the virtual machine and memory usage as described in the following table.

Text Box	Option
Virtual Machine Name	Enter a name for the virtual machine being created; for example, VM-0A1B2C3D4. **NOTE:** The default name is the name of the source machine.
Version	Specify the version of VMware Workstation for the virtual machine. You can choose from: □ VMware Workstation 7.0 □ VMware Workstation 8.0 □ VMware Workstation 9.0.
Memory	Specify the memory for the virtual machine. □ Click **Use the same amount of RAM as the source machine** to specify that the RAM configuration is the same as the source machine. Or, □ 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.

4. Click **Export**.

Performing a Continuous (Virtual Standby) VMware Workstation Export

You can choose to perform a continuous export for VMware Workstation Export. Complete the steps in this procedure to perform a continuous export.

To perform a continuous (virtual standby) VMware Workstation export

1. In the AppAssure 5 Core Console, click the Machines tab.
2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **VMware Workstation Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **Continuous (Virtual Standby)** and then click **Next**.

   The VM Export - Select Recovery Point dialog box displays.
5. Select a recovery point to export and then click **Next.**

   The Virtual Standby Recovery Point to VMware Workstation/Server dialog box displays.
6. Enter the parameters for accessing the virtual machine as described in the following table.

Text Box	Option
Target Path	Specify the path of the local folder or network share on which to create the virtual machine. **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.
User name	Enter the logon credentials for the virtual machine. □ If you specified a network share path, you need to 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	Enter the logon credentials for the virtual machine. □ If you specified a network share path, you need to 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.

7. In the Export Volumes pane, select the volumes to export; for example, C:\ and D:\.

8. In the Options pane, enter the information for the virtual machine and memory usage as described in the following table.

Text Box	Option
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.
Memory	Specify the memory for the virtual machine. □ Click **Use the same amount of RAM as the source machine** to specify that the RAM configuration is the same as the source machine. Or, □ 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.

9. Click **Perform initial ad-hoc export** to test the export of the data.
10. Click **Save**.

Exporting Windows Data using Hyper-V Export

In AppAssure 5, 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

You can choose to perform a one-time export for Hyper-V. Complete the steps in this procedure to perform a one-time export.

To perform a one-time Hyper-V export

1. In the AppAssure 5 Core Console, click the Machines tab.
2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **Hyper-V Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **One-time export**.
5. Click **Next**.

   The Hyper-V Export - Select Recovery Point dialog box displays.
6. Select a recovery point to export and then click **Next.**

   The Hyper-V dialog box displays.

Defining One-Time Settings for Performing a Hyper-V Export

Complete the steps in this procedure to define the settings for performing a one-time 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 as described in the following table.

Text Box	Description
Hyper-V 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.
VM Machine Location	Enter the path for the virtual machine; for example, D:\export. It is used to identify the location of the virtual machine. **NOTE:** Specify the virtual machine location for both local and remote Hyper-V servers. The path should be a valid local path for the Hyper-V server. Non-existent directories are automatically created. You should not attempt to create them manually. Export to shared folders, for example, \\data\share is not permitted.

3. In the Export Volumes tab, select which volumes to export; for example, C:\.
4. Select the Options tab and then enter the name for the virtual machine in the **Virtual Machine Name** text box.

   The name you enter appears in the list of virtual machines in the Hyper-V Manager console.
5. Do one of the following:
   - Click **Use the same amount of RAM as the source machine** to identify that the RAM use is identical between the virtual and source machines.
   - Click **Use a specific amount of RAM** to specify how much memory the virtual machine should have after the export; for example, 4096 MB.
6. Click **Export**.

Performing a Continuous (Virtual Standby) Hyper-V Export

You can choose to perform a continuous export for Hyper-V. Complete the steps in this procedure to perform a continuous export.

To perform a one-time Hyper-V export

1. In the AppAssure 5 Core Console, click the Machines tab.
2. In the list of protected machines, select the machine or cluster with the recovery point that you want to export.
3. In the Actions drop-down menu for that machine, click **Export**, and then select **Hyper-V Export**.

   The Select Export Type dialog box displays.
4. In the Select Export Type dialog box, click **Continuous (Virtual Standby)**.
5. Click **Next**.

   The Hyper-V dialog box displays.
6. Click the **Use local machine** option perform the Hyper-V export to a local machine with the Hyper-V role assigned.
7. 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 as described in the following table.

Text Box	Description
Hyper-V 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.
VM Machine Location	Enter the path for the virtual machine; for example, D:\export. It is used to identify the location of the virtual machine. **NOTE:** Specify the virtual machine location for both local and remote Hyper-V servers. The path should be a valid local path for the Hyper-V server. Non-existent directories are automatically created. You should not attempt to create them manually. Export to shared folders, for example, \\data\share is not permitted.

8. In the Export Volumes tab, select which volumes to export; for example, C:\.

9. Select the Options tab and then enter the name for the virtual machine in the **Virtual Machine Name** text box.

   The name you enter appears in the list of virtual machines in the Hyper-V Manager console.

10. Do one of the following:

    - Click **Use the same amount of RAM as the source machine** to identify that the RAM use is identical between the virtual and source machines.
    - Click **Use a specific amount of RAM** to specify how much memory the virtual machine should have after the export; for example, 4096 MB.

11. Click **Perform initial ad-hoc export** to test the export of the data.

12. Click **Save**.

Performing a Rollback

In AppAssure 5, a rollback is the process of restoring the volumes on a machine from recovery points.

To perform a rollback

1. In the AppAssure 5 Core Console, do one of the following:

   - Click the Machines tab, and then do the following:
     i. In the list of protected machines, select the check box next to the machine you want to export.
     ii. In the Actions drop-down menu for that machine, click **Rollback.**
     iii. In the Rollback - Select Recovery Point dialog box, select a recovery point to export and click **Next.**
   - Or, in the left navigation area, select the machine you want to roll back, which launches the Summary tab for that machine.
     iv. Click the Recovery Points tab, and then select a recovery point from the list.
     v. Expand the details for that recovery point, and then click **Rollback.**

   The Rollback - Choose Destination dialog box displays.

2. Edit the rollback options as described in the following table.

Text Box	Description
Protected Machine	Specify the original agent machine as the destination for the rollback. Source refers to the agent from which the recovery point being used was created.
Recovery Console Instance	Enter the user name and password to restore the recovery point to any machine that booted in URC mode.

3. Click **Load Volumes.**

   The Volume Mapping dialog box displays.

   NOTE: The Core console does not automatically map Linux volumes. To locate a Linux volume, browse to the volume that you want to roll back.

4. Select the volumes that you want to roll back.
5. Using the **Destination** drop-down lists, select the destination volume to which the selected volume should roll back.
6. Select the options as described in the following table.

Option	Description
Live Recovery	When selected, the rollback for Windows volumes happens immediately. Selected by default. **NOTE:** The Live Recovery option is not available for Linux volumes.
Force Dismount	When selected, it forces the dismount of any mounted recovery points prior to performing the rollback. Selected by default.

7. Click **Rollback**.

   The system begins the process of rolling back to the selected recovery point.

Performing a Rollback for a Linux Machine by Using the Command Line

A rollback is the process of restoring the volumes on a machine from recovery points. In AppAssure 5, you can perform a rollback for volumes on your protected Linux machines using the command line `aamount` utility.

Note

When performing this procedure, do not attempt to mount recovery points to the /tmp folder, which contains the aavdisk files.

Caution

You should not attempt to perform a rollback on the system or root (/) volume.

Note

Rollback functionality is also supported for your protected machines within the AppAssure 5 Core Console. See "Performing a Rollback" on page 153 for more information.

To perform a rollback for a volume on a Linux machine

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 hostname of your AppAssure Core server.

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

   A list displays showing the machines protected by this AppAssure server. 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: Note that you can also enter the machine ID number in this command instead of the line item number.

   A list displays that shows the base and incremental recovery points for that machine. This list includes a line item number, date/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. Enter the following command to select a recovery point for rollback:

   ```
   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: You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, you would use the 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 rollback the `23` recovery point volume `b` to the volume that was mounted to the directory `/mnt/data`, the command would be:

   ```
   r2 23 b /mnt/data
   ```

   NOTE: It is possible to rollback to /, but only when performing a Bare Metal Restore while booted with a Live CD. For more information, see "Launching a Bare Metal Restore for Linux" on page 177.

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

   Once the rollback proceeds, a series of messages will display to notify you of the status.

8. Upon a successful rollback, the `aamount` utility will automatically mount and re-attach the kernel module to the rolled back volume if the target was previously protected and mounted. If not, you will need to mount the rollback volume to the local disk and then should verify that the files are restored. (for example, you can use the `sudo mount` command and then the `ls` command.)

Note that you should not unmount a protected Linux volume by hand. In the event you need to do this, 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 would need to be in a form similar to this example:
`/dev/sda1`.

Understanding Bare Metal Restore

Servers, when operating as expected, perform the tasks they are configured to do. It is only when they fail that things change. When a catastrophic event occurs, rendering a server inoperable, immediate steps are needed to restore the full functionality of that machine.

AppAssure 5 provides the ability to perform a bare metal restore (BMR) for your Windows or Linux machines. BMR is a process that restores the full software configuration for a specific system. It uses the term "bare metal" because the restore operation recovers not only the data from the server, but also reformats the hard drive and reinstalls the operating system and all software applications. To perform a BMR, you specify a recovery point from a protected machine, and roll back to the designated physical or virtual machine. Other circumstances in which you may choose to perform a bare metal restore include hardware upgrade or server replacement.

Windows 8, 8.1 and Windows Server 2012, 2012 R2 operating systems that are booted from FAT32 EFI partitions are not available for protection or recovery, nor are Resilient File System (ReFS) volumes.

Bare metal restore of Storage Spaces disks configuration (a feature of Windows 8.1) is also not supported in this release.

Only supported Linux operating systems are available for protection or recovery. This includes Ubuntu, Red Hat Enterprise Linux, CentOS, and SUSE Linux Enterprise Server (SLES). For details, see the AppAssure 5 Deployment Guide.

Performing a BMR is possible for physical or virtual machines. As an added benefit, AppAssure 5 allows you to perform a BMR whether the hardware is similar or dissimilar. Performing a BMR on AppAssure 5 separates the operating system from a specific platform, providing portability.

Examples of performing a BMR for similar hardware include replacing the hard drive of the existing system, or swapping out the failed server with an identical machine.

Examples of performing a BMR for dissimilar hardware include restoring a failed system with a server produced by a different manufacturer or with a different configuration. This process encompasses creating a boot CD image, burning the image to disk, starting up the target server from the boot image, connecting to the recovery console instance, mapping volumes, initiating the recovery, and then monitoring the process. Once the bare metal restore is complete, you can continue with the task of loading the operating system and the software applications on the restored server, followed by establishing unique settings required for your configuration.

Bare metal restore is used not only in disaster recovery scenarios, but also to migrate data when upgrading or replacing servers.

While BMR is supported for virtual machines, it is also worth noting that it is easier to perform a Virtual Export for a VM than it is to perform a BMR on a physical machine. For more information on performing a VM export for virtual machines, see the appropriate procedure for the supported VM.

- For more information on performing a VM export using ESXi, see "Exporting Windows Data using ESXi Export" on page 144.
- For more information on performing a VM export using VMware Workstation, see "Exporting Windows Data using VMware Workstation Export" on page 147.
- For more information on performing a VM export using Hyper-V, see "Exporting Windows Data using Hyper-V Export" on page 150.

To perform a BMR on a Windows machine, refer to the roadmap specific to Windows, including the prerequisites. For more information, see "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

To perform a BMR on a Linux machine, refer to the roadmap specific to Linux, including prerequisites. In addition to performing a BMR using the command line aamount utility, you can now perform a BMR from within the Core Console UI. The roadmap takes both approaches into account. For more information, see "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171.

Roadmap for Performing a Bare Metal Restore for a Windows Machine

To perform a bare metal restore for Windows machines, perform the following tasks.

- **Manage a Windows boot image**. This boot CD ISO image will be used to start up the destination drive, from which you can access the Universal Recovery Console to communicate with backups on the Core. See "Managing a Windows Boot Image" on page 159.
  - If you require physical media to start up the destination machine, you will need to **transfer the boot CD ISO image to media**. See "Transferring the Boot CD ISO Image to Media" on page 163.

- In all cases, you will need to **load the boot image into the destination server and start the server** from the boot image. See "Loading the Boot CD and Starting the Target Machine" on page 163.

- **Launch a Bare Metal Restore for Windows**. Once the destination machine is started from the boot CD, you can launch the BMR. See "Launching a Bare Metal Restore for Windows" on page 164.
  - You will need to **initiate rollback from a recovery point on the Core**. See "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165.
  - You will need to **map the volumes**. See "Mapping Volumes for a Bare Metal Restore" on page 166.
  - If restoring to dissimilar hardware, you will need to **inject drivers for hardware devices** that were not in the previous configuration but are included in the system replacing the server. For more information, see "Injecting Drivers to Your Target Server" on page 167.
- **Verifying a Bare Metal Restore**. After starting the bare metal restore, you can verify and monitor your progress. See "Verifying a Bare Metal Restore" on page 168.
  - You can **monitor the progress of your restore**. See "Viewing the Recovery Progress" on page 169.
  - Once completed, you can **start the restored server**. See "Starting a Restored Target Server" on page 169
  - **Troubleshoot the BMR process**. See "Troubleshooting Connections to the Universal Recovery Console" on page 170 and "Repairing Startup Problems" on page 170.

Prerequisites for Performing a Bare Metal Restore for a Windows Machine

Before you can begin the process of performing a bare metal restore for a Windows machine, you must ensure that the following conditions and criteria exist:

- **Backups of the machine you want to restore.** You must have a functioning AppAssure 5 Core containing recovery points of the protected server you want to restore
- **Hardware to restore (new or old, similar or dissimilar).** The target machine must meet the installation requirements for an agent; for details, see the AppAssure 5 Deployment Guide.
- **Image media and software.** You must have a blank CD or DVD and disk burning software, or software to create an ISO image. If managing machines remotely using virtual network computing software such as UltraVNC, then you must have VNC Viewer.
- **Compatible storage drivers and network adapter drivers.** If restoring to dissimilar hardware, then you must have Windows 7 PE (32-bit) compatible storage drivers and network adapter drivers for the target machine, including RAID, AHCI, and chipset drivers for the target operating system, as appropriate.

- **Storage space and partitions, as appropriate.** 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.
- **Compatible partitions.** Windows 8 and Windows Server 2012 operating systems that are booted from FAT32 EFI partitions are not available for protection or recovery, nor are Resilient File System (ReFS) volumes. If the protected Windows machine contains non-supported partitions such as FAT32, these partitions will not be transferred to the restored machine. For details, see the AppAssure 5 Deployment Guide.

Managing a Windows Boot Image

A bare metal restore for Windows requires a boot image referred to as the boot CD, which you create by defining parameters in the AppAssure 5 Core Console. This image is tailored to your specific needs. You will use the image to start the destination Windows machine. Based on the specifics of your environment you may need to transfer this image to physical media such as a CD or DVD. You must then virtually or physically load the boot image, and start the Windows server from the boot image.

This process is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

To manage a Windows boot image, you can perform the following tasks:

- "Creating a Boot CD ISO Image for Windows" on page 159
- "Defining Boot CD ISO Image Parameters" on page 160
- "Transferring the Boot CD ISO Image to Media" on page 163
- "Loading the Boot CD and Starting the Target Machine" on page 163

Creating a Boot CD ISO Image for Windows

The first step when performing a bare metal restore (BMR) for a Windows machine is to create the boot CD file in the AppAssure 5 Core Console. This is a bootable ISO image which contains the AppAssure 5 Universal Recovery Console (URC) interface, an environment that is used to restore the system drive or the entire server directly from the AppAssure 5 Core.

The boot CD 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 the recovery point originated, then you must include storage controller and other drivers in the boot CD. For information about injecting those drivers in the boot CD, see "Injecting Drivers in a Boot CD" on page 161.


Note

The International Organization for Standardization (ISO) is an international body of representatives from various national organizations that sets file system standards. The ISO 9660 is a file system standard that is used for optical disk media for the exchange of data and supports various operating systems, for example, 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.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Windows Boot Image" on page 159.

To create a boot CD ISO image

1. From the AppAssure 5 Core Console where the server you need to restore is protected, select the Core and then click the Tools tab.
2. Click **Boot CDs**.
3. Select **Actions**, and then click **Create Boot CD**.

   The Create Boot CD dialog box displays. Use the following procedures to complete the dialog box.

Defining Boot CD ISO Image Parameters

Once you open the Create Boot CD dialog box, there are several parameters that may be required. Based on the specifics of your situation, perform the following tasks as required to define properties for a boot CD ISO image to use for a bare metal restore.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Creating a Boot CD ISO Image for Windows" on page 159.

Naming the Boot CD File and Setting the Path

Complete the following step to name the boot CD file and set the path where the ISO image is stored.

To name the boot CD file and set the path

- In the Create Boot CD dialog box, in Output Options, in the Output path text box, enter the path where you want to store the boot CD ISO image on the Core server.

  If the shared drive 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, the backslash (only as a path delimiter), 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

Complete the following steps to create the 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 four 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 require 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

The boot CD image requires storage drivers to recognize the drives of the server, and network adapter drivers in order to communicate with the AppAssure 5 Core over the network.

A generic set of Windows 7 PE 32-bit storage controller and network adapter drivers are included automatically when you generate a boot CD for Windows. This will satisfy the requirements of newer Dell systems. Systems from other manufacturers or older Dell systems may require you to inject storage controller or network adapter drivers when creating the boot CD.

When creating the boot CD, driver injection is used to facilitate the operability between the recovery console, network adapter, and storage on the target server.

Data restored from the recovery point includes drivers for the hardware previously in place. If performing a bare metal restore to dissimilar hardware, then you must also inject storage controller drivers into the operating system being restored using the URC after the data has been restored to the drive, This allows the restored operating system to boot using the new set of hardware. Once the OS is booted after the restore, you can then download and install any additional drivers needed by the OS to interact with its new hardware.

For more information, see "Injecting Drivers to Your Target Server" on page 167.

Complete the following steps to inject storage controller and network adapter drivers in a boot CD.

To inject drivers in a boot CD

1. Download the drivers from the manufacturer's Web site for the server and unpack them.
2. Compress each driver into a .zip file using an appropriate compression utility (for example, WinZip).
3. In the Create Boot CD dialog box, in the Drivers pane, click **Add a Driver**.
4. Navigate through the filing system to locate the compressed driver file, select the file, and then click **Open**.

   The injected drivers appear highlighted in the Drivers pane.
5. Repeat Step 3 and Step 4, as appropriate, until all drivers have been injected.

Creating the Boot CD ISO Image

Complete the following step to create the boot CD ISO image.

To create a boot CD ISO image

- After you have named the boot CD file 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 and saved with the filename you provided.

Viewing the ISO Image Creation Progress

Complete the following step to view the progress of the creation of the ISO image.

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 will appear on the Boot CDs page, accessible from the Tools menu.

Accessing the ISO Image

Complete the following step to access the ISO image.

To access 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.

Transferring the Boot CD ISO Image to Media

When you create the boot CD file, it is stored as an ISO image in the path you specified. You must be able to mount this image as a drive on the server on which you are performing a bare metal restore.

You can burn the boot CD ISO image onto compact disc (CD) or digital video disk (DVD) media accessible at system startup.

When you start the machine from the boot CD, the Universal Recovery Console launches automatically.

If performing a BMR on a virtual machine, this step is not required. Simply load the ISO image in a drive and edit settings for that VM to start from that drive.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Windows Boot Image" on page 159.

Loading the Boot CD and Starting the Target Machine

After you create the boot CD image, you need to boot the target server with the newly created boot CD.

If you created the boot CD using DHCP, you must capture the IP address and password.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Windows Boot Image" on page 159.

To load a boot CD and start the target machine

1. Navigate to the new server and load the boot CD image from the appropriate location. Specify that the server will start from the boot CD image.
2. Start the machine, which loads the following:
   - Windows 7 PE

- AppAssure 5 Agent software

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

NOTE: A new temporary password is generated each time the machine is started with the boot CD. Write down the IP address displayed in the Network Adapters Settings pane and the authentication password displayed in the Authentication pane. You will need this information later during the data recovery process to log back on to the console.

3. If you want to change the IP address, select it and click **Change**.

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

Once started with the boot CD, this machine is ready for the user to connect to it from the Core to begin the bare metal restore process.

Launching a Bare Metal Restore for Windows

Before launching a bare metal restore (BMR) for a Windows machine, certain conditions are required.

To restore a recovery point saved on the Core, you must have the appropriate hardware in place. For more information, see "Prerequisites for Performing a Bare Metal Restore for a Windows Machine" on page 158.

The BMR destination Windows machine must be started using the boot CD image. For more information, see "Managing a Windows Boot Image" on page 159.

The first step is to select the appropriate recovery point, then initiate the rollback to the hardware by specifying the IP address and temporary password you obtained from the Universal Recovery Console.

You must then map the drives and start the rollback.

The recovery point includes drivers from the previous hardware. If restoring to dissimilar hardware, then you must inject storage controller drivers into the operating system being restored using the URC after the data has been restored to the drive, This allows the restored operating system to boot using the new set of hardware. Once the OS is booted after the restore, you can then download and install any additional drivers needed by the OS to interact with its new hardware.

To launch a BMR from the AppAssure 5 Core Console, perform the following tasks.

- "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165
- "Mapping Volumes for a Bare Metal Restore" on page 166
- "Injecting Drivers to Your Target Server" on page 167

This process is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

Selecting a Recovery Point and Initiating Rollback for BMR

Once the Universal Recovery Console is accessible on the machine on which you want to perform a BMR, you must select the recovery point that you want to restore. Navigate to the Core Console to select which recovery point you want to load, and designate the recovery console as the destination for the restored data.


Note

This step is required to perform BMR on all Windows machines and optional to perform BMR on Linux machines.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Launching a Bare Metal Restore for Windows" on page 164.

If performing a BMR for a Linux machine from the Core Console, then this task is also a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171 It is part of the process for "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179.

Complete the steps in this procedure to select a recovery point on the Core to roll back to the physical or virtual BMR target machine.

To select a recovery point and initiate a rollback for BMR

1. Navigate to the AppAssure 5 Core Console and, in the list of protected machines, click the name of the protected server you want to restore to bare metal.

   The Summary tab for the selected machine appears.

2. Click the Recovery Points tab.

3. In the list of recovery points, click the right angle bracket > symbol to expand the recovery point that you want to restore.

4. In the expanded details for that recovery point, from the Actions menu, click **Rollback**.

   The Rollback - Choose Destination dialog box appears.

5. Select **Recovery Console Instance**.

   The authentication fields become accessible.

6. Enter the information about the machine to which you want to connect as described in the following table, and then click **Load Volumes**.

Text Box	Description
Host	The IP address of the machine to which you want to restore. This is identical to the IP address you wrote down from the URC Console.
Password	The specific password to connect to the selected server. This is identical to the Current Password shown in the URC Console.

If the connection information you entered matches the URC console, and if the Core and the target server can identify each other properly on the network, then the volumes for the selected recovery point are loaded, and the RollbackURC dialog box appears, In this case, your next step is to map volumes.

NOTE: If the protected Windows machine contains non-supported partitions such as FAT32 or ReFS, these partitions will not be transferred to the restored machine. Bare metal restore of Storage Spaces disks configuration (a feature of Windows 8.1) is also not supported in this release. For details, see the AppAssure 5 Deployment Guide.

Mapping Volumes for a Bare Metal Restore

Once connected to the Universal Recovery Console, you will need to map volumes between those listed in the recovery point and volumes existing on the target hardware to perform the restore.

AppAssure 5 attempts to automatically map volumes. If you accept the default mapping, then the disk on the destination machine is cleaned and re-partitioned and **any previously existing** data is deleted. The alignment is performed in the order the volumes are listed **in the recovery point,** and the volumes are allocated to the disks appropriately according to size, and so on. **Assuming there is enough space on the target drive, no partitioning is required when using automatic disk alignment.** A disk can be used by multiple volumes. If you manually map the drives, note that you cannot use the same disk twice.

For manual mapping, you must have the new machine correctly formatted already before restoring it. The destination machine **must have a separate partition for each volume in the recovery point, including the system reserved volume**. For more information, see "Launching a Bare Metal Restore for Windows" on page 164.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 153. It is part of the process for "Launching a Bare Metal Restore for Windows" on page 164.

If performing a BMR for a Linux machine from the Core Console, then this task is also a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Launching a Bare Metal Restore for Linux" on page 177.

Complete the steps in this procedure to map a volume.

To map volumes for a bare metal restore

1. If you want to map volumes automatically, do the following. If you want to map volumes manually, proceed to Step 2
   a. In the RollbackURC dialog box, 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 listed beneath it and are selected.

      NOTE: Typically for a BMR, you should restore, at minimum, the system reserved volume and the system volume (usually, but not always, the C:\ volume).
   c. Optionally, if you do not wish to restore a listed volume, clear the option under Source volume. At least one volume must be selected to perform the BMR.
   d. If the destination disk that is automatically mapped is the correct target volume, select **Destination Disk** and ensure that all appropriate volumes are selected.
   e. Click **Rollback**, and then proceed to Step 3.
2. If you want to map volumes manually, do the following:
   a. In the RollbackURC dialog box, select the Manually Map Volumes tab.

      NOTE: If no volumes exist on the drive of the machine on which you are performing a BMR, you will not be able to see this tab or manually map volumes.
   b. In the Volume Mapping area, under Source Volume, verify that the source volume is selected, and that the appropriate volumes are listed beneath it and are selected.
   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 **Begin Rollback**.

If you select **Begin Rollback**, 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.

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 once you reboot the system following the restore process.

If you are unsure which drivers are required by your target server, 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

Your target server automatically contains some generic Windows 7 PE 32-bit drivers which will work for some systems.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Launching a Bare Metal Restore for Windows" on page 164.

Complete the following task to inject drivers to your target server.

To inject drivers to your target server

1. Download the drivers from the manufacturer's Web site for the server and unpack them.
2. Compress each driver into a .zip file using an appropriate compression utility (for example, WinZip) and copy it to the target server.
3. In the Universal Recovery Console, click **Driver Injection.**
4. Navigate through the filing system to locate the compressed driver file 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.
6. Repeat Step 3 through Step 5, as appropriate, until all drivers have been injected.

Verifying a Bare Metal Restore

Once you perform a bare metal restore, you can verify the progress of the restore. When the action is completed successfully, you can start the restored server. Some troubleshooting steps are included if you encounter difficulties connecting to the Universal Recovery Console to complete the restore, and to repair startup problems with the restored machine.

You can perform the following tasks:

- "Viewing the Recovery Progress" on page 169
- "Starting a Restored Target Server" on page 169
- "Troubleshooting Connections to the Universal Recovery Console" on page 170
- "Repairing Startup Problems" on page 170

This process is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

Viewing the Recovery Progress

Complete the steps in this procedure to view the recovery progress of a rollback initiated from the AppAssure 5 Core Console.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Verifying a Bare Metal Restore" on page 168.

To view the recovery progress

1. After you initiate the rollback process, the Active Task dialog box appears, showing that the rollback action initiated.

   NOTE: This does not indicate a successful completion of the rollback.

2. Optionally, to monitor the rollback task, from the Active Task dialog box, click **Open Monitor Window**.

   NOTE: From the Monitor Open Task window, you can view the status of the recovery, as well as the start and end times.

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

Starting a Restored Target Server

Complete the steps in this procedure to start the restored target server.

Before starting the restored target server, you should verify that the recovery was successful. For more information, see "Viewing the Recovery Progress" on page 169.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Verifying a Bare Metal Restore" on page 168.

To start a restored target server

1. Navigate back to the target server, and verify that the AppAssure Universal Recovery Console is active.
2. Eject the boot CD (or disconnect physical media with the boot CD image) from the restored server.
3. In the Universal Recovery Console, from the Console tab, click **Reboot**.
4. Specify to start the operating system normally.
5. Log on to the machine. The system should be restored to its state captured in the recovery point.

Troubleshooting Connections to the Universal Recovery Console

The following are troubleshooting steps for connecting to the boot CD image as part of the process for "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Verifying a Bare Metal Restore" on page 168.

If an error displays indicating that the Core could not connect to the remote server, then any of several possible causes are likely.

- Verify that the IP address and Current Password displayed in the URC is identical to the information you entered in the Recovery Console Instance dialog box.
- To reach the server on which to restore data, the Core must be able to identify the server on the network. To determine if this is possible, you can open a command prompt on the Core and ping the IP address of the target BMR server. You can also open a command prompt on the target server and ping the IP address of the AppAssure 5 Core.
- Verify that the network adapter settings are compatible between Core and target BMR server.

Repairing Startup Problems

Complete the steps in this procedure to repair startup problems. Keep in mind that if you restored to dissimilar hardware, you must have injected storage controller, RAID, AHCI, chipset and other drivers if they are not already on the boot CD. These drivers make it possible for the operating system to operate all devices on your target server successfully. For more information, see "Injecting Drivers to Your Target Server" on page 167.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Verifying a Bare Metal Restore" on page 168.

Complete the following procedure to repair startup problems on your target server.

To repair startup problems

1. 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**.

Roadmap for Performing a Bare Metal Restore on Linux Machines

In AppAssure 5 you can perform a Bare Metal Restore (BMR) for a Linux machine, including rollback of the system volume. When you restore a Linux machine, you will roll back to the boot volume recovery point. BMR functionality is supported using the command line aamount utility and from within the Core Console UI.

To perform a bare metal restore for Linux machines, perform the following tasks.

- **Manage a Linux boot image**. This Linux Live DVD boot ISO image is used to start up the destination drive, from which you can access the Universal Recovery Console to communicate with backups on the Core. See "Managing a Linux Boot Image" on page 173.
  - If you require physical media to start up the destination Linux machine, you will need to **transfer the ISO image to media**. See "Transferring the Live DVD ISO Image to Media" on page 174.
  - In all cases, you will need to **load the boot image into the destination server and start the server** from the boot image. See "Loading the Live DVD and Starting the Target Machine" on page 174.
- **Manage Partitions**. You may need to create or mount partitions before performing a BMR on a Linux machine. See"Managing Linux Partitions" on page 175.
  - The Linux system on which you are performing a BMR must have the same partitions as the source volumes in the recovery point. You may need to **create additional partitions on the target system**, if required. See "Creating Partitions on the Destination Drive" on page 175.
  - **Mount partitions.** If performing a BMR from the Core Console, you must first mount partitions. See "Mounting Partitions from the Command Line" on page 177. Steps to mount partitions are included in the process to perform a BMR from the command line. See "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179.
- **Launch a Bare Metal Restore for Linux**. Once the destination machine is started from the Live DVD boot image, you can launch the BMR. The tasks required depend on whether you will perform this from the AppAssure user interface or from the command line using the aamount utility. See "Launching a Bare Metal Restore for Linux" on page 177.
  - If using the Core Console, you will need to **initiate rollback from a recovery point on the Core**. See "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165.
  - If using the Core Console, you will need to **map the volumes** from the UI. See "Mapping Volumes for a Bare Metal Restore" on page 166.
  - Optionally, if restoring from the command line, you can **start the screen utility** to enhance your ability to scroll and see commands in the terminal console. For more information, see "Starting the Screen Utility" on page 178.
  - **If using aamount, all tasks will be performed at the command line**. For more information, see "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179.

- **Verifying a Bare Metal Restore**. After starting the bare metal restore, you can verify and monitor your progress. See "Verifying the Bare Metal Restore from the Command Line" on page 182.
  - You can **monitor the progress of your restore**. See "Viewing the Recovery Progress" on page 169.
  - Once completed, you can **start the restored server**. See "Starting a Restored Target Server" on page 169.
  - **Troubleshoot the BMR process**. See "Troubleshooting Connections to the Universal Recovery Console" on page 170 and "Repairing Startup Problems" on page 170.

Prerequisites for Performing a Bare Metal Restore for a Linux Machine

Before you can begin the process of performing a bare metal restore for a Linux machine, you must ensure that the following conditions and criteria exist:

- **Backups of the machine you want to restore.** You must have a functioning AppAssure 5 Core containing recovery points of the protected server you want to restore.
- **Hardware to restore (new or old, similar or dissimilar).** The target machine must meet the installation requirements for an agent; for details, see the AppAssure 5 Deployment Guide.
- **Live DVD boot image**. Obtain the Linux Live DVD ISO image, which includes a bootable version of Linux. Download it from the license portal at https://licenseportal.com. If you have any issues downloading the Live DVD, contact Dell AppAssure support.
- **Image media and software.** If using physical media, you must have a blank CD or DVD and disk burning software, or software to create an ISO image.
- **Compatible storage drivers and network adapter drivers.** If restoring to dissimilar hardware, then you must have compatible storage drivers and network adapter drivers for the target machine, including RAID, AHCI, and chipset drivers for the target operating system, as appropriate.
- **Storage space and partitions, as appropriate.** 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.
- **Rollback path.** 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.

Managing a Linux Boot Image

A bare metal restore for Linux requires a Live DVD boot image, which you download from the license portal. You will use this image to start the destination Linux machine. Based on the specifics of your environment you may need to transfer this image to physical media such as a CD or DVD. You must then virtually or physically load the boot image, and start the Linux server from the boot image.


Note

The Live DVD was previously known as the Live CD.

You can perform the following tasks:

- "Downloading a Boot ISO Image for Linux" on page 173
- "Transferring the Live DVD ISO Image to Media" on page 174
- "Loading the Live DVD and Starting the Target Machine" on page 174

Managing a Linux boot image is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171.

Downloading a Boot ISO Image for Linux

The first step when performing a bare metal restore (BMR) for a Linux machine is to download the Linux Live DVD ISO image from the license portal. The Live DVD functions with all Linux file systems supported by AppAssure 5, and includes a bootable version of Linux, a screen utility, and the AppAssure Universal Recovery Console (URC) interface. The AppAssure 5 Universal Recovery Console is an environment that is used to restore the system drive or the entire server directly from the AppAssure 5 Core.


Note

The International Organization for Standardization (ISO) is an international body of representatives from various national organizations that sets file system standards. The ISO 9660 is a file system standard that is used for optical disk media for the exchange of data and supports various operating systems. 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.

You must download the Live DVD ISO image that matches your version of AppAssure 5. The current version of Live DVD is available from the license portal at https://licenseportal.com. If you need a different version, contact Dell AppAssure Support.


Note

For more information about the license portal, see "Managing AppAssure 5 Licenses", on the AppAssure 5 Technical Documentation page at: http://docs.appassure.com/display/AA50D/AppAssure+5+Technical+Documentation.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Linux Boot Image" on page 173.

Complete the steps in this procedure to download the Live DVD ISO image.

To download a Boot ISO image for Linux

1. Log into the License Portal at https://licenseportal.com.
2. Access the **Downloads** area.
3. Scroll down to Linux Based Applications and, from the Linux Live CD section, click **Download**.
4. Save the Live DVD ISO image. If you are restoring a virtual machine, you can save it to a network location, and set the VM to start up from a CD or DVD drive associated with the ISO image.
5. If restoring from a physical machine, burn the Boot CD ISO image onto a compact disc (CD) or digital video disk (DVD) from which the target machine can be started. For more information, see "Transferring the Live DVD ISO Image to Media" on page 174.

Transferring the Live DVD ISO Image to Media

When you download the Linux Live DVD file, it is stored as an ISO image in the path you specified. You must be able to boot the target Linux machine from the Live DVD image.

You can burn the boot CD ISO image onto compact disc (CD) or digital video disk (DVD) media.

When you start the machine from the Live DVD, the Universal Recovery Console launches automatically.

If performing a BMR on a virtual machine, this step is not required. Simply load the ISO image in a drive and edit settings for that VM to start from that drive.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Linux Boot Image" on page 173.

Loading the Live DVD and Starting the Target Machine

After you obtain the Live DVD ISO image, you need to start the Linux machine from the newly created Live DVD.

This task is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157. It is part of the process for "Managing a Linux Boot Image" on page 173.

To load a Live DVD and start the target machine

1. Navigate to the new server and load the Live DVD image from the appropriate location. Specify that the server will start from the Live DVD image.
2. Start the machine.

   An AppAssure splash screen displays and a terminal window opens, displaying the IP address and authentication password for the machine.

   NOTE: A new temporary password is generated each time the machine is started with the Live DVD image.
3. Write down the IP address and the authentication password displayed on the introduction screen. You will need this information later during the data recovery process to log back on to the console.

Once the target Linux machine is started with the Live DVD, this machine is ready for the user to connect to it from the Core to begin the bare metal restore process. You can perform this process using any one of two methods:

- Launching a restore from the AppAssure 5 Core Console. For more information, see "Launching a Bare Metal Restore for Linux" on page 177.
- Launching a Restore from the command Line using the aamount utility. For more information, see "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179.

Managing Linux Partitions

When performing a BMR, the destination drive onto which you will be restoring data must have the same partitions as in the recovery point you are restoring. You may need to create partitions to meet this requirement.

You can launch the restore from the command line using the aamount utility, or you can launch the restore from the AppAssure 5 Core Console. If restoring using the user interface, you must first mount the partitions.

You can perform the following tasks:

- "Creating Partitions on the Destination Drive" on page 175
- "Mounting Partitions from the Command Line" on page 177

Managing Linux partitions is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171.

Creating Partitions on the Destination Drive

Often, when performing a BMR, the destination drive is a new volume that may consist of a single partition. The drive on the destination machine must have the same partition table as in the recovery point, including the size of the volumes. If the destination drive does not contain the same partitions, you must create them before performing the bare metal restore. Use the fdisk utility to create partitions on the destination drive equal to the partitions on the source drive.

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Managing Linux Partitions" on page 175.

To create partitions on the destination drive

1. Optionally, you can start the Screen utility. Once started, it remains active until you reboot the machine.

   NOTE: For more information, see "Starting the Screen Utility" on page 178.

2. From the command line, enter the following command and then press **Enter** to change privileges to run as administrator and then list existing disk partitions:

   ```
   sudo fdisk -l
   ```

   A list of all volumes appears.

   This example assumes the volume you want to partition is /dev/sda. If your volume is different (for example, for older drives, you may see /dev/hda), change commands accordingly.

3. To create a new boot partition, enter the following command and then press **Enter**:

   ```
   sudo fdisk /dev/sda
   ```

4. To create a new boot partition, enter the following command and then press **Enter**:

   ```
   n
   ```

5. To create a new primary partition, enter the following command and then press **Enter**:

   ```
   p
   ```

6. To specify partition number, enter the partition number and then press **Enter**. For example, to specify partition 1, type 1 and then press **Enter**.

7. To use the first sector, 2048, press **Enter**.

8. Allocate an appropriate amount to the boot partition by entering the plus sign and the allocation amount and then press **Enter**.

   For example, to allocate 500 M for the boot partition, type the following and then press **Enter**:

   ```
   +500M
   ```

9. To toggle a bootable flag for the boot partition (to make the partition bootable), type the following command and then press **Enter**:

   ```
   a
   ```

10. To assign a bootable flag for the appropriate partition, type the number of the partition and then press **Enter**. For example, to assign a bootable flag for partition 1, type 1 and then press **Enter**.

11. To save all changes in the fdisk utility, type the following command and then press **Enter**:

    ```
    w
    ```

Mounting Partitions from the Command Line

If performing a BMR using the AppAssure 5 Core Console, you must first mount the appropriate partitions on the destination machine. Perform this from the command line in the Universal Recovery Console.

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Managing Linux Partitions" on page 175.

Complete the steps in this procedure to mount partitions on the Linux machine before performing a rollback.

To mount partitions from the command line

1. From the command line, enter the following command and then press **Enter** to change privileges to run as administrator and then list existing disk partitions:

   ```
   sudo fdisk -l
   ```

   A list of all volumes appears.

2. Mount all partitions you will need to perform the BMR to the mount directory. These must match the volumes that are in the recovery point. For example, if the volume you want to mount is called sda1, and the mount directory is mnt, then type the following command and then press **Enter**:

   ```
   mount /dev/sda1 /mnt
   ```

3. Repeat Step 2 as necessary until you have mounted all required volumes.

Once volumes are mounted, you can perform a rollback to the destination Linux machine from the AppAssure 5 Core Console. See "Launching a Bare Metal Restore for Linux" on page 177.

Launching a Bare Metal Restore for Linux

Before launching a bare metal restore (BMR) for a Linux machine, certain conditions are required.

To restore a recovery point saved on the Core, you must have the appropriate hardware in place. For more information, see "Prerequisites for Performing a Bare Metal Restore for a Linux Machine" on page 172.

The BMR destination Linux machine must be started using the Live DVD boot image. For more information, see "Managing a Linux Boot Image" on page 173.

The number of volumes on the Linux machine to be restored must match the number of volumes in the recovery point. You must also decide whether to restore from the AppAssure 5 Core Console, or from the command line using aamount. For more information, see "Managing Linux Partitions" on page 175.

If restoring from the Core Console UI, the first step in launching a BMR is to select the appropriate recovery point, then initiate the rollback to the hardware by specifying the IP address and temporary password you obtained from the Universal Recovery Console. You must then map the drives and start the rollback.

To launch a BMR from the AppAssure 5 Core Console, perform the following tasks.

- "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165
- "Mapping Volumes for a Bare Metal Restore" on page 166
- "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165

If restoring from the command line using the aamount utility, then you must first set appropriate privileges, mount volumes, execute aamount, obtain information about the Core from the list of machines, connect to the core, obtain a list of recovery points, select the recovery point you want to roll back onto bare metal, and launch the rollback.

Optionally, you may want to start the Screen utility.

To launch a BMR from the command line, perform the following tasks.

- "Starting the Screen Utility" on page 178
- "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179

This process is a step in "Roadmap for Performing a Bare Metal Restore for a Windows Machine" on page 157.

Starting the Screen Utility

Included on the Live DVD is Screen, a utility which is available when you boot from the Live DVD into the Universal Recovery Console. Screen allows users to manage multiple shells simultaneously over a single Secure Shell (SSH) session or console window. This allows you to perform one task in a terminal window (such as verify mounted volumes) and, while that is running, open or switch to another shell instance to perform another task (such as to run the aamount utility).

The Screen utility also has its own scroll-back buffer, which enables you to scroll the screen to view larger amounts of data, such as a list of recovery points.


Note

This utility is provided for convenience; use of the Screen utility is optional.

Before you can use it, you must start the Screen utility from the Live DVD using the procedure below. Once you reboot the machine, the utility closes.

To start the screen utility

1. Using the Live DVD file, start the Linux machine.

   An AppAssure splash screen displays and a terminal window opens.

2. At the command prompt, type `screen` and press **Enter** to start the screen utility.

Launching a Bare Metal Restore for a Linux Machine using the Command Line

Once the Live DVD ISO image is accessible on the machine on which you want to perform a BMR, and the number and size of volumes matches between the target machine and the recovery point you want to restore to bare metal, then you can launch a restore from the command line using the `aamount` utility.

If you want to perform a BMR restore using the AppAssure 5 Core Console UI, see "Selecting a Recovery Point and Initiating Rollback for BMR" on page 165.


Note

When performing this procedure, do not attempt to mount recovery points to the /tmp folder, which contains the aavdisk files.

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Launching a Bare Metal Restore for a Linux Machine using the Command Line" on page 179.

Complete the steps in this procedure to select a recovery point on the Core to roll back to the physical or virtual BMR target machine.

To perform a bare metal restore for a Linux machine using the command line

1. On the Linux machine that is the destination of the bare metal restore, from the command prompt of the Universal Recovery Console (URC), optionally, you can start the Screen utility. Once started, it remains active until you reboot the machine.

   NOTE: For more information, see "Starting the Screen Utility" on page 178.

2. Type the following command to see if the appropriate partitions are mounted, and then press **Enter**:

   ```
   df
   ```

3. If the volumes you need are mounted, skip to Step 5. If not, then to mount the volumes, type the following command and then press **Enter**:

   ```
   mount <volume> <folder>
   ```

   For example, if the volume path is dev/sda1 and the folder you want to mount to is mnt, then type the following and then press **Enter**:

   ```
   mount /dev/sda1 /mnt
   ```

4. To run the AppAssure `aamount` utility as root, type the following command and then press **Enter**:

   ```
   sudo aamount
   ```

5. To list the protected machines, type the following command and then press **Enter**:

   ```
   lm
   ```

6. When prompted, enter the connection information for the AppAssure 5 Core as described in the following table, pressing **Enter** after each required command:

Text Box	Description	Required
AppAssure Core IP address or hostname	The IP address or hostname of the AppAssure 5 Core.	Yes
Domain	The domain of the AppAssure 5 Core. This is optional.	No
User	The user name for an administrative user on the Core	Yes
Password	The password used to connect the administrative user to the Core.	Yes

A list displays showing the machines protected by the AppAssure Core. It lists the machines found by line item number, the host display name or IP address, and an ID number for the machine.

```
Found 2 machine(s):

          Host/Address                                                ID
  1        10.10.61.16                  0f001b7c-eabb-4409-a1e5-4d94adb656ce
  2        10.10.61.15                  293cc667-44b4-48ab-91d8-44bc74252a4f
```

Figure 24. List of protected machines

7. To list the recovery points for the machine that you want to restore, type the list recovery points command using the following syntax and then press **Enter**:

```
lr <machine_line_item_number>
```


Note

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

A list displays the base and incremental recovery points for that machine. This list includes:

- A line item number
- Date and time stamp
- A lettered list of volumes within the recovery point
- Location of the volume

- Size of the recovery point
- An ID number for the volume that includes a sequence number at the end, which identifies the recovery point

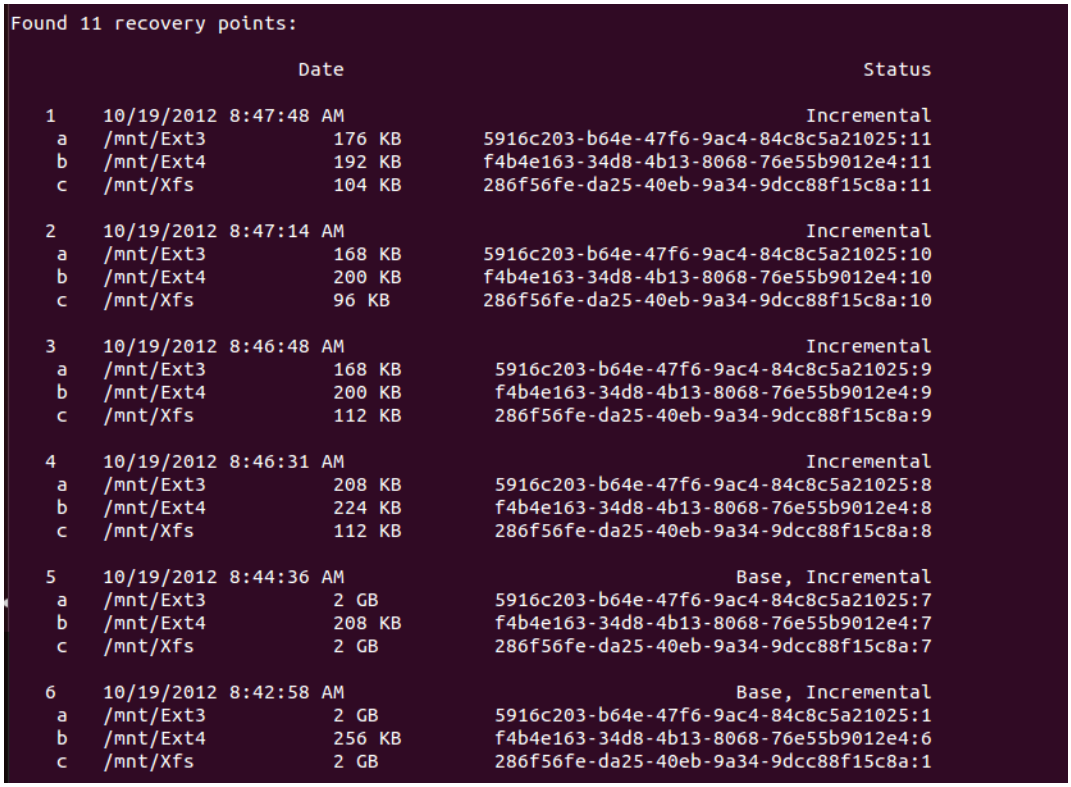

Figure 25. List of recovery points

8. To select the recovery point for rollback, enter the following command and then press **Enter**:

   r <recovery_point_ID_number> <path>

Caution

You must ensure that the system volume is not mounted.

Note

If you started the machine from the Live DVD, then the system volume is not mounted.

This command 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.

You can also specify a line number in the command instead of the recovery point ID number to identify the recovery point. In that case, use the agent/machine line number (from the `lm` output), followed by the recovery point line number and volume letter (from the lettered list of volumes within the recovery point), followed by the path. For example:


Note

```
r <machine_line_item_number> <base_image_recovery_point_line_number>
<volume_letter> <path>
```

For example, type:

```
r 1 24 a /dev/sda1
```

In this command, <path> is the file descriptor for the actual volume.

9. When prompted to proceed, enter **y** for Yes and then press **Enter**.

   After the rollback begins, a series of messages will display that notify you of the rollback completion status.

   NOTE: If you receive an exception message, the details regarding that exception can be found in the aamount.log file. The aamount.log file is located in /var/log/appassure.

10. Upon a successful rollback, exit aamount by typing **exit** and then press **Enter**.

11. Your next step is to verify the restore. For more information, see "Verifying the Bare Metal Restore from the Command Line" on page 182.

Verifying the Bare Metal Restore from the Command Line

Dell recommends performing the following steps to verify a bare metal restore completed from the command line.

- "Performing a File System Check on the Restored Volume" on page 182
- "Creating Bootable Partitions on the Restored Linux Machine using the Command Line" on page 183

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171.

Performing a File System Check on the Restored Volume

Once you execute a bare metal restore from the command line, you should perform a file system check on the restored volume to ensure the data restored from the recovery point was not corrupted.

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Verifying the Bare Metal Restore from the Command Line" on page 182.

Perform the task below to perform a file system check on the restored volume.

To perform a file system check on the restored volume

1. From the command line in the Universal Recovery Console of the Linux machine you have restored, to verify whether the appropriate partitions are mounted, type the following command and then press **Enter**:

   ```
   df
   ```

2. If the restored volume is not mounted, then skip to Step 3. If the restored volume is mounted, unmount it by typing the following command and then pressing **Enter**:

   ```
   umount <volume>
   ```

3. Run a file system check on the restored volumes by typing the following command and then press Enter:

   ```
   fsck <volume>
   ```

   If the fsck returns clean, the file system is verified.

4. Mount the appropriate volumes once again by typing the following command in format **`mount <volume> <folder>`**, and then press **Enter**.

   For example, if the volume path is prod/sda1 and the folder you want to mount to is mnt, then type the following and then press **Enter**:

   ```
   mount /dev/sda1 /mnt
   ```

Creating Bootable Partitions on the Restored Linux Machine using the Command Line

Once you complete a clean file system check on the restored volume, you must create bootable partitions.

GNU Grand Unified Bootloader (GRUB) is a boot loader that allows administrators to configure which operating system or specific kernel configuration is used to start the system. After a BMR, the configuration file for GRUB must be modified so that the machine uses the appropriate universally unique identifier (UUID) for the root volume. Before this step you must mount the root and boot volumes, and check the UUIDs for each. This ensures that you can boot from the partition.

This task is a step in "Roadmap for Performing a Bare Metal Restore on Linux Machines" on page 171. It is part of the process for "Verifying the Bare Metal Restore from the Command Line" on page 182.

Perform the task below to create bootable partitions using the command line.

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

1. From the command line in the Universal Recovery Console of the Linux machine you have restored, attach to all devices using the bsctl utility with the following command as root:

   ```
   sudo bsctl --attach-to-device /<restored volume path>
   ```

   For example, if the volume path is dev/sda1 and the folder you want to mount to is mnt, then type the following and then press **Enter**:

   ```
   sudo bsctl --attach-to-device /dev/sda1 mnt
   ```

   NOTE: Repeat this step for each restored volume.

2. You must mount the root volume first and then the boot volume. Mount each restored volume by using the following commands:

   a. To mount the root volume, type the following command and then press **Enter**:

      ```
      mount /<restored volume[root]> /mnt
      ```

      For example, if /dev/sda2 is the root volume, then type **`mount /dev/sda2 /mnt`** and then press **Enter**.

   b. To mount the boot volume, type the following command and then press **Enter**:

      ```
      mount /<restored volume[boot]> /mnt/boot
      ```

      For example, if /dev/sda1 is the boot volume, then type **`mount /dev/sda1 /mnt/boot`** and then press **Enter**.

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

3. If the volume size is increasing — that is, if the destination volume on the new Linux machine is larger than the volume was in the recovery point — then you must delete any existing bitmap data files, and then recreate them as described in Step 4 through Step 7,

   If the source and target volumes are the same size, proceed to Step 8 to reset the bitmap store.

   For both situations, you will then need to map them as described in Step 9.

This is a critical step prior to mapping volumes. If you map a volume and then delete the file manually, you could corrupt the volume.

---

4. If the volume size is increasing, then delete the existing data store by typing the following command and then press **Enter**:

   ```
   rm -rf <mount point>/.blksnap/data
   ```

   For example, if your restored volume was mounted to /mnt/sda1, then type the following command and then press **Enter**:

   ```
   rm -rf /mnt/sda1/.blksnap/data
   ```

5. Now you must delete the existing bitmap store by typing the following command and then press **Enter**:

   ```
   rm -rf <mount point>/.blksnap/bitmap
   ```

   For example, if your restored volume was mounted to /mnt/sda1, then type the following command and then press **Enter**:

   ```
   rm -rf /mnt/sda1/.blksnap/bitmap
   ```

6. If you deleted the existing data store and bitmap store, then recreate the data store by typing the following command and then press **Enter**:

   ```
   sudo bsctl --create-data-store <restored root volume path>
   ```

   For example, type the following command and then press **Enter**:

   ```
   sudo bsctl --create-data-store /dev/sda1
   ```

7. Repeat this for the bitmap store by typing the following command and then press **Enter**:

   ```
   sudo bsctl --create-bitmap-store <restored root volume path>
   ```

   For example, type the following command and then press **Enter**:

   ```
   sudo bsctl --create-bitmap-store /dev/sda1
   ```

8. f the source and target volumes are the same size, reset the bitmap store by typing the following command and then pressing **Enter**:

   ```
   sudo bsctl --reset-bitmap-store <restored volume path>
   ```

   For example, type the following command and then press **Enter**:

   ```
   sudo bsctl --reset-bitmap-store /dev/sda1
   ```

9. For all situations, map snapshot metadata for each restored volume by using the following command and then press **Enter**:

   ```
   sudo bsctl --map-bitmap-store <restored volume path>
   ```

   For example, type the following command and then press **Enter**:

   ```
   sudo bsctl --map-bitmap-store /dev/sda1
   ```

10. Verify that the devices are mapped by typing **bsctl -l** and pressing **Enter**.

11. Obtain the Universally Unique Identifier (UUID) of the new volumes by using the `blkid` command. Type the following and then press **Enter**:

    ```
    blkid [volume]
    ```

    NOTE: You can also use the `ls -l /dev/disk/by-uuid` command.

12. Obtain the UUID of `mount /etc/fstab` and compare it to the UUIDs for the root (for Ubuntu and CentOS) and boot (for CentOS and RHEL) volumes by typing the following command and then press **Enter**:

    ```
    less /mnt/etc/fstab
    ```

13. Obtain the UUID of `mount /etc/mtab` and compare it to the UUIDs for the root (for Ubuntu and CentOS) and boot (for CentOS and RHEL) volumes by typing the following command and then press **Enter**:

    ```
    less /mnt/etc/mtab
    ```

14. If performing a BMR on a brand new disk on the destination machine, comment out the swap partition in fstab in your root volume.

15. Modifying fstab and mtab paths should occur on the restored volume, not the Live CD. There is no need to modify paths on the Live CD. Prepare for the installation of Grand Unified Bootloader (GRUB) by typing the following commands. Following each command, press **Enter**:

    ```
    mount --bind /dev /mnt/dev
    ```

    ```
    mount --bind /proc /mnt/proc
    ```

16. Locate the `grub.conf` file in your mounted volume, and open it using a text editor.

    The location of grub.conf differs depending on your OS version and the version of GRUB installed. The most likely locations include **<root path>/boot/grub/grub.conf**, **<root path>/boot/grub/grub.cfg** or **<root path>/etc/grub.conf**.

17. In grub.conf, locate all lines containing "root=<root device uuid>" and replace it with the correct UUID for the root volume. If not, update each instance so that root=<root device uuid>. You can also use the root device path. As in the examples above, if the root device path is /dev/sda2, then change all instances to root=/dev/sda2.

18. Remove all "rd_LVM_LV=" entries in the grub.conf file, save the file, and exit the text editor.

19. Change root directory by typing the following command and then press **Enter**:

    ```
    chroot /mnt /bin/bash
    ```

20. Install GRUB by typing the following command and then press **Enter**:

    ```
    grub-install/dev/sda
    ```

    NOTE: If installing on SUSE, when installing GRUB, no parameters are required. For example, the command to install GRUB on SUSE is simply grub-install and then press **Enter**.

21. Remove the Live DVD disk from the CD-ROM or DVD drive and restart the Linux machine.

Viewing Events and Alerts

Complete the steps in this procedure to view events and alerts for a machine.

To view events and alerts for a machine

1. Do one of the following:

- To view events for all machines managed by a core, in the Core Console, click the Events tab.
- To view events for a specific machine, in the AppAssure 5 Core Console, navigate to the machine for which you want to view events, and click the Events tab.

The Events tab displays a log of all events. The contents of the Events tab is divided into two sections, Tasks and Alerts, for you to view details about task and alert events respectively.

2. In the upper right area of the Events tab, you can control how active and completed events are displayed by doing the following:
   - To show only active events, ensure only **Active** is selected.
   - To show only completed events, ensure that only **Completed** is selected.
   - To show both active and completed events, ensure that both **Active** and **Completed** are selected.
3. If you want to remove all alerts from the page, click **Dismiss All.**

This page is intentionally left blank.

5 Protecting Server Clusters

This chapter describes how to protect information on Microsoft SQL Server or Exchange Server clusters using AppAssure 5. It includes the following topics:

About Server Cluster Protection in AppAssure 5

In AppAssure 5, server cluster protection is associated with the AppAssure Agents installed on individual cluster nodes (that is, individual machines in the cluster) and the AppAssure 5 Core, which protects those agents, all as if they were one composite machine.

You can easily configure an AppAssure 5 Core to protect and manage a cluster. In the Core Console, a cluster is organized as a separate entity, which acts as a 'container' to include the related nodes. For example, in the left navigation area, the Core is listed at the top of the navigation tree, and then clusters are listed under the Core and contain the associated individual nodes (on which the AppAssure Agents are installed).

At the Core and cluster levels, you can view information about the cluster, such as the list of related nodes and shared volumes. A cluster appears in the Core Console on the Machines tab, and you toggle the view (using Show/Hide) to view the nodes included in the cluster. At the cluster level, you can also view corresponding Exchange and SQL cluster metadata for the nodes in the cluster. You can specify settings for the entire cluster and the shared volumes in that cluster, or you can navigate to an individual node (machine) in the cluster to configure settings just for that node and the associated local volumes.

Supported Applications and Cluster Types

To protect your cluster properly, you must have installed the AppAssure 5 Agent on each of the machines or nodes in the cluster. AppAssure 5 supports the application versions and cluster configurations listed in the following table.

Application	Application Version and Related Cluster Configuration	Windows Failover Cluster
Microsoft Exchange	2007 Single Copy Cluster (SCC) 2007 Cluster Continuous Replication (CCR)	2003, 2008, 2008 R2
	2010 Database Availability Group (DAG)	2008, 2008 R2
Microsoft SQL	2005, 2008, 2008 R2 Single Copy Cluster (SCC)	2003, 2008, 2008 R2
	2012 Single Copy Cluster (SCC)	2008, 2008 R2, 2012

The supported disk types include:

- GUID partition table (GPT) disks greater than 2 TB
- Basic disks

The supported mount types include:

- Shared drives that are connected as drive letters (for example, D:)
- Simple dynamic volumes on a single physical disk (not striped, mirrored, or spanned volumes)
- Shared drives that are connected as mount points

Protecting a Cluster

This topic describes how to add a cluster for protection in AppAssure 5. When you add a cluster to protection, you need to specify the host name or IP address of the cluster, the cluster application, or one of the cluster nodes or machines that includes the AppAssure 5 Agent.

Note

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

For information about setting up repositories, see "About Repositories" on page 36.

To protect a cluster

1. Do one of the following:
   - In the Core Console, navigate to the Home tab, and then click the **Protect Cluster** button.
   - In the Core Console, on the Machines tab, click **Actions**, and then click **Protect Cluster**.
2. In the Connect to Cluster dialog box, enter the following information, and then click **Connect**.

Text Box	Description
Host	The host name or IP address of the cluster, the cluster application, or one of the cluster nodes that you wish to protect. **NOTE:** If you use the IP address of one of the nodes, this node needs to have an AppAssure 5 Agent installed and started.
Port	The port number on the machine on which the AppAssure 5 Core communicates with the Agent.
User name	The user name of the domain administrator used to connect to this machine: for example, domain_name\administrator or administrator@domain_name.com **NOTE:** The domain name is mandatory. You cannot connect to the cluster using the local administrator user name.
Password	The password used to connect to this machine.

3. In the Protect Cluster dialog box, select a repository for this cluster.
4. To protect the cluster based on default settings, select the nodes for default protection, and click **Protect**.

   NOTE: The default settings ensure that all volumes are protected with a schedule of every 60 minutes.
5. To enter custom settings for the cluster (for example, to customize the protection schedule for the shared volumes), do the following:
   a. Click **settings**.

b. In the Volumes dialog box, select the volume(s) to protect, and then click **Edit**.

c. In the Protection Schedule dialog box, select one of the schedule options for protecting your data as described in the following table.

Text Box	Description
Interval	You can choose from: □ **Weekday**. To protect data on a specific interval, select **Interval**, and then: □ To customize when to protect data during peak times, you can specify a start time, end time, and an interval. □ To protect data during off-peak times, select the **Protect during off-peak times** check box, and then select an interval for protection. □ **Weekends**. To protect data during weekends as well, select the **Protect during weekends** check box, and then select an interval.
Daily	To protect data on a daily basis, select the **Daily** option, and then for Protection Time, select a time to start protecting data.
No Protection	To remove protection from this volume, select the **No Protection** option.

6. When you have made all necessary changes, click **Save**.
7. To enter custom settings for a node in the cluster, select a node, and then click the **Settings** link next to the node.
   - Repeat Step 5 to edit the protection schedule.

   For more information on customizing nodes, see "Protecting Nodes in a Cluster" on page 192.
8. In the Protect Cluster dialog box, click **Protect**.

Protecting Nodes in a Cluster

This topic describes how to protect the data on a cluster node or machine that has an AppAssure 5 Agent installed. When you add protection, you need to select a node from the list of available nodes as well as specify the host name and the user name and password of the domain administrator.

To protect nodes in a cluster

1. Once you have added a cluster, navigate to the Home tab, and then under Protected Machines, select the cluster.

   The Summary tab for the selected cluster displays.
2. Click the Protected Nodes tab, and then from the **Actions** menu, select **Protect Cluster Node**.

3. In the Protect Cluster Node dialog box, select or enter as appropriate the following information, and then click **Connect** to add the machine or node.

Text Box	Description
Host	A drop-down list of nodes in the cluster available for protection.
Port	The port number on which the AppAssure 5 Core communicates with the Agent on the node.
User name	The user name of the domain administrator used to connect to this node; for example, example_domain\administrator or administrator@example_domain.com.
Password	The password used to connect to this machine.

4. Click **Protect** to start protecting this machine with default protection settings.

   NOTE: The default settings ensure that all volumes on the machine are protected with a schedule of every 60 minutes.

5. To enter custom settings for this machine, (for example, to change the Display name, add encryption, or customize the protection schedule), click **Show Advanced Options**.

6. Edit the following settings as needed, as described in the following table.

Text Box	Description
Display Name	Enter a new name for the machine to be displayed in the Core Console.
Repository	Select the repository on the AppAssure 5 Core in which the data from this machine should be stored.
Encryption	Specify whether encryption should be applied to the data for every volume on this machine to be stored in the repository. **NOTE:** The encryption settings for a repository are defined under the Configuration tab in the AppAssure 5 Core Console.
Schedule	Select one of the following options. □ Protect all volumes with default schedule □ Protect specific volumes with custom schedule. Then, under Volumes, select a volume and click **Edit**. For more information on setting custom intervals, see Step 5 in "Protecting a Cluster" on page 191.

Process of Modifying Cluster Node Settings

Once you have added protection for cluster nodes, you can easily modify basic configuration settings for those machines or nodes (for example, display name, host name, and so on), protection settings (for example, changing the protection schedule for local volumes on the machine, adding or removing volumes, and pausing protection), and more.

To modify cluster node settings, you must perform the following tasks:

1. Do one of the following.
   - Navigate to the cluster that contains the node you want to modify, click the Machines tab, and select the machine or node that you want to modify.
   - Or, from the Navigation pane, under the Cluster heading, select the machine or node you want to modify.
2. To modify and view configuration settings, see "Viewing and Modifying Configuration Settings" on page 98.
3. To configure notification groups for system events, see "Configuring Notification Groups for System Events" on page 100.
4. To customize retention policy settings, see "Customizing Retention Policy Settings" on page 105.
5. To modify the protection schedule, see "Modifying Protection Schedules" on page 108.
6. To modify transfer settings, see "Modifying Transfer Settings" on page 110.

Roadmap for Configuring Cluster Settings

The roadmap for configuring cluster settings involves performing the following tasks:

- **Modify cluster settings.** For more information about modifying cluster settings, see "Modifying Cluster Settings" on page 194.
- **Configure cluster event notifications.** For more information about configuring cluster event notifications, see "Configuring Cluster Event Notifications" on page 195.
- **Modify the cluster retention policy.** For more information about modifying the cluster retention policy, see "Modifying the Cluster Retention Policy" on page 196.
- **Modify the cluster protection schedules.** For more information about modifying the cluster protection schedules, see "Modifying Cluster Protection Schedules" on page 197.
- **Modify the cluster transfer settings.** For more information about modifying cluster transfer settings, see "Modifying Cluster Transfer Settings" on page 197.

Modifying Cluster Settings

Once you have added a cluster, you can easily modify basic settings (for example, display name), protection settings (for example, protection schedules, adding or removing volumes, and pausing protection), and more.

To modify cluster settings

1. Do one of the following:

- In the Core Console, click the Machines tab, and then select the cluster you wish to modify.
- Or, in the left navigation area, select the cluster you wish to modify.

2. Click the Configuration tab.

   The Settings page displays.

3. Click **Edit** to modify the settings on this page for the cluster as described in the following table.

Text Box	Description
Display Name	Enter a display name for the cluster. The name for this cluster displays in the AppAssure 5 Core Console. By default, this is the host name for the cluster. You can change this to something more descriptive, if needed.
Host Name	This setting represents the host name for the cluster. It is listed here for informational purposes only and cannot be modified.
Repository	Enter the Core repository associated with the cluster. **NOTE:** If snapshots have already been taken for this cluster, this setting is listed here for informational purposes only and cannot be modified.
Encryption Key	Edit and select an encryption key if necessary. This specifies whether encryption should be applied to the data for every volume on this cluster to be stored in the repository.

Configuring Cluster Event Notifications

You can configure how system events are reported for your cluster by creating notification groups. These events could be system alerts or errors. Complete the steps in this procedure to configure notification groups for events.

To configure cluster event notifications

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster you wish to modify.
   - Or, in the left navigation area, select the cluster you wish to modify.
2. Click the Configuration tab, and then click **Events**.
3. Select one of the options described in the following table.

Option	Description
Use Core alert settings	This adopts the settings used by the associated core: □ Click **Apply** and then perform Step 5.
Use Custom alert settings	This lets you configure custom settings: □ Proceed to Step 4.

4. If you selected Custom alert settings, do the following:

   a. Click **Add Group** to add a new notification group for sending a list of system events.

      The Add Notification Group dialog box opens.

   b. Add the notification options as described in the following table.

Text Box	Description
Name	Enter a name for the notification group.
Description	Enter a description for the notification group.
Enable Events	Select the events for notification, for example, Clusters. You can also choose to select by type: □ Error □ Warning □ Info **NOTE:** When you choose to select by type, by default, the appropriate events are automatically enabled. For example, if you choose Warning, the Attachability, Jobs, Licensing, Archive, CoreService, Export, Protection, Replication, and Rollback events are enabled.
Notification Options	Select the method to specify how to handle notifications You can choose from the following options: □ **Notify by Email**. Specify the email addresses to which to send the events in the To, CC, and BCC text boxes. □ **Notify by Windows Event log**. The Windows Event log controls the notification. □ **Notify by syslogd.** Specify the host name and port to which to send the events.

   c. Click **OK** to save your changes, and then click **Apply**.

5. To edit an existing notification group, next to a notification group in the list, click **Edit**.

   The Edit Notification Group dialog box displays for you to edit the settings.

Modifying the Cluster Retention Policy

The retention policy for a cluster specifies how long the recovery points for the shared volumes in the cluster are stored in the repository. Retention policies are used to retain backup snapshots for longer periods of time and to help with management of these backup snapshots. The retention policy is enforced by a rollup process that helps in aging and deleting old backups.

To modify the cluster retention policy

1. Do one of the following:

- In the Core Console, click the Machines tab, and then select the cluster you wish to modify.
- Or, in the left navigation area, select the cluster you wish to modify.

2. Click the Configuration tab, and then click **Retention Policy**.
3. Select one of the options in the following table.

Option	Description
Use Core default retention policy	This adopts the settings used by the associated core. □ Click **Apply**.
Use Custom retention policy	This lets you configure custom settings. □ Proceed to Step NOTE:.

NOTE: If you selected Custom alert settings, follow the instructions for setting a custom retention policy as described in "Customizing Retention Policy Settings" on page 105, beginning with Step 4.

Modifying Cluster Protection Schedules

In AppAssure 5, you can modify the protection schedules only if your cluster has shared volumes.

To modify cluster protection schedules

1. In the Core Console, click the Machines tab, and then select the cluster you wish to modify.
2. Follow the instructions for modifying the protection settings as described in "Modifying Protection Schedules" on page 108.

Modifying Cluster Transfer Settings

In AppAssure 5, you can modify the settings to manage the data transfer processes for a protected cluster.

You can modify cluster transfer settings only if your cluster has shared volumes.

There are three types of transfers in AppAssure 5:

- **Snapshots**. Backs up the data on your protected cluster.
- **VM Export**. Creates a virtual machine with all of the backup information and parameters as specified by the schedule defined for protecting the cluster.

- **Rollback**. Restores backup information for a protected cluster.

To modify cluster transfer settings

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster you wish to modify.
   - Or, in the left navigation area, select the cluster you wish to modify.
2. Click the Configuration tab, and then click **Transfer Settings**.
3. Modify the protection settings as described in "Modifying Protection Schedules" on page 108, beginning with Step 2.

Converting a Protected Cluster Node to an Agent

In AppAssure 5, you can convert a protected cluster node to an AppAssure Agent so that it is still managed by the Core, but it is no longer part of the cluster. This is helpful, for example, if you need to remove the cluster node from the cluster but still keep it protected.

To convert a protected cluster node to an agent

1. Do one of the following.
   - In the Core Console, click the Machines tab, and select the cluster that contains the machine you wish to convert. Then, click the Machines tab for the cluster.
   - Or, from the left navigation area, select the cluster that contains the machine you want to convert, and click the Machines tab.
2. Select the machine to convert, click the **Actions** drop-down menu at the top of the Machines tab, and click **Convert to Agent**.
3. To add the machine back to the cluster, select the machine, and then click the Summary tab, the **Actions** menu, and **Convert to Node.**

Viewing Server Cluster Information

Complete the steps in the following procedures to view summary, event, alert information, and so on for server clusters.

Viewing Cluster System Information

Complete the steps in this procedure to view detailed system information about a cluster.

To view cluster system information

1. Do one of the following:

- In the Core Console, click the Machines tab, and then select the cluster you wish to view.
- Or, in the left navigation area, select the cluster you wish to view.

2. Click the Tools tab.

   The system information page displays to show system details about the cluster such as name, included nodes with associated state and Windows versions, network interface information, and volume capacity information.

Viewing Cluster Events and Alerts

Complete the steps in this procedure to view events and alerts for a cluster.

For information about viewing events and alerts for an individual machine or node in a cluster, see "Viewing Events and Alerts" on page 186.

To view cluster events and alerts

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster you wish to view.
   - Or, in the left navigation area, under Clusters, select the cluster you wish to view.
2. Click the Events tab, which opens to show a log of all events for current tasks as well as any alerts for the cluster.
3. To filter the list of events, you can select or clear the **Active**, **Complete**, or **Failed** check boxes as appropriate.
4. In the Alerts table, click **Dismiss All** to dismiss all of the alerts in the list.

Viewing Summary Information

Complete the steps in this procedure to view summary information about a cluster including information about the associated quorum for the cluster.

To view summary information

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster you wish to view.
   - Or, in the left navigation area, under Clusters, select the cluster you wish to view.
2. On the Summary tab, you can view such information as the cluster name, cluster type, quorum type (if applicable), and the quorum path (if applicable). This tab also shows at-a-glance information about the volumes in this cluster, including size and protection schedule.
3. To refresh this information to the most current, click the **Actions** drop-down menu, and click **Refresh Metadata**.

For information about viewing summary and status information for an individual machine or node in the cluster, see "Viewing Machine Status and Other Details" on page 120.

Working with Cluster Recovery Points

A recovery point, also referred to as a snapshot, is a point-in-time copy of the folders and files for the shared volumes in a cluster, which are stored in the repository. Recovery points are used to recover protected machines or to mount to a local file system. In AppAssure 5, you can view the lists of recovery points in the repository. Complete the steps in the following procedure to review recovery points.


Note

If you are protecting data from a DAG or CCR server cluster, the associated recovery points do not appear at the cluster level. They are only visible at the node or machine level.

For information about viewing recovery points for individual machines in a cluster, see "Viewing Recovery Points" on page 133.

To work with cluster recovery points

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster for which you wish to view recovery points.
   - Or, in the left navigation area, under Clusters, select the cluster for which you wish to view recovery points.
2. Click the Recovery Points tab.
3. To view detailed information about a specific recovery point, click the right angle bracket > symbol next to a recovery point in the list to expand the view.
4. For information about the operations you can perform on the recovery points, see "Viewing a Specific Recovery Point" on page 134
5. Select a recovery point to mount.

   For information about how to mount a recovery point, see "Mounting a Recovery Point for a Windows Machine" on page 136.
6. To delete recovery points, see "Removing Recovery Points" on page 139.

Managing Snapshots for a Cluster

In AppAssure 5, you can manage snapshots by forcing a snapshot or by pausing current snapshots. Forcing a snapshot lets you force a data transfer for the currently protected cluster. When you force a snapshot, the transfer starts immediately or will be added to the queue. Only the data that has changed from a previous recovery point transfers. If there is no previous recovery point, all data (the base image) on the protected volumes is transferred. When you pause a snapshot, you temporarily stop all transfers of data from the current machine.

For information about forcing snapshots for the individual machines in a cluster, see "Forcing a Snapshot" on page 141. For information about pausing and resuming snapshots for the individual machines in a cluster, see "Pausing and Resuming Protection" on page 98.

Forcing a Snapshot for a Cluster

Complete the steps in this procedure to force a snapshot for a cluster.

To force a snapshot for a cluster

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster for which you wish to view recovery points.
   - Or, in the left navigation area, under Clusters, select the cluster for which you wish to view recovery points.
2. On the Summary tab, click the **Actions** drop-down menu, and then click **Force Snapshot**.

Pausing and Resuming Cluster Snapshots

Complete the steps in this procedure to pause and resume a snapshot for a cluster.

To pause and resume cluster snapshots

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster for which you wish to view recovery points.
   - Or, in the left navigation area, under Clusters, select the cluster for which you wish to view recovery points.
2. On the Summary tab, click the **Actions** drop-down menu, and then click **Pause Snapshots**.

3. In the Pause Protection dialog box, select one of the options described in the following table.

Option	Description
Pause until resumed	Pauses the snapshot until you manually resume protection. □ To resume protection, click the **Actions** menu and then click **Resume**.
Pause for	Lets you specify an amount of time in days, hours, and minutes to pause snapshots.

Dismounting Local Recovery Points

Complete the steps in this procedure to dismount recovery points that are mounted locally.

To dismount local recovery points

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster for which you wish to dismount recovery points.
   - Or, in the left navigation area, select the cluster for which you wish to dismount recovery points.
2. On the Tools tab, under the Tools menu, click **Mounts**.
3. In the list of local mounts, do one of the following:
   - To dismount a single local mount, locate and select the mount for the recovery point you want to dismount, and then click **Dismount**.
   - To dismount all local mounts, click the **Dismount All** button.

Performing a Rollback for Clusters and Cluster Nodes

A rollback is the process of restoring the volumes on a machine from recovery points. For a server cluster, you perform a rollback at the node, or machine, level. This section provides guidelines for performing a rollback for cluster volumes.

Performing a Rollback for CCR (Exchange) and DAG Clusters

Complete the steps in this procedure to perform a rollback for CCR (Exchange) and DAG clusters.

To perform a rollback for CCR (Exchange) and DAG clusters

1. Turnoff all nodes except one.

2. Perform a rollback using the standard AppAssure procedure for the machine as described in "Performing a Rollback" on page 153 and "Performing a Rollback for a Linux Machine by Using the Command Line" on page 154.
3. When the rollback is finished, mount all databases for the cluster volumes.
4. Turn on all other nodes.
5. For Exchange, navigate to the Exchange Management Console, and, for each database, perform the **Update Database Copy** operation.

Performing a Rollback for SCC (Exchange, SQL) Clusters

Complete the steps in this procedure to perform a rollback for SCC (Exchange, SQL) clusters.

To perform a rollback for SCC (Exchange, SQL) clusters

1. Turnoff all nodes except one.
2. Perform a rollback using the standard AppAssure procedure for the machine as described in "Performing a Rollback" on page 153 and "Performing a Rollback for a Linux Machine by Using the Command Line" on page 154.
3. After the rollback is finished, mount all databases from the cluster volumes.
4. Turn on all other nodes one-by-one.

   NOTE: You do not need to roll back the quorum disk. It can be regenerated automatically or by using cluster service functionality.

Replicating Cluster Data

When you are replicating data for a cluster, you configure replication at the machine level for the individual machines in that cluster. You can also configure replication to replicate the recovery points for shared volumes; for example, if you have five agents that you want to replicate from source to target.

For more information and instructions on replicating data, see "Replicating Agent Data on a Machine" on page 118.

Removing a Cluster from Protection

Complete the steps in the following procedure to remove a cluster from protection.

To remove a cluster from protection

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster you wish to remove.
   - Or, in the left navigation area, select the cluster you wish to remove to view the Summary tab.
2. Click the **Actions** drop-down menu, and then click **Remove Machine**.
3. Select one of the following options.

Option	Description
Keep Recovery Points	To keep all currently stored recovery points for this cluster.
Remove Recovery Points	To remove all currently stored recovery points for this cluster from the repository.

Removing Cluster Nodes from Protection

Complete the steps in the following procedures to remove cluster nodes from protection.

If you just want to remove a node from the cluster, see "Converting a Protected Cluster Node to an Agent" on page 198.

To remove a cluster node from protection

1. Do one of the following.
   - In the Core Console, click the Machines tab, and then select the cluster that contains the node you wish to remove. On the Machines tab for the cluster, select the node you want to remove.
   - Or, in the left navigation area, under the related cluster, select the node you want to remove.
2. Click the Actions drop-down menu and then click **Remove Machine**.
3. Select one of the options described in the following table.

Option	Description
Keep Recovery Points	To keep all currently stored recovery points for this machine or node.
Remove Recovery Points	To remove all currently stored recovery points for this machine or node from the repository.

Removing All Nodes in a Cluster from Protection

Complete the steps in this procedure to remove all nodes in a cluster from protection.

To remove all nodes in a cluster from protection

1. Do one of the following.
   - In the Core Console, click the Machines tab, and select the cluster that contains the nodes you wish to remove. Then, click the Machines tab for the cluster.
   - Or, from the left navigation area, select the cluster that contains the nodes you want to remove, and then click the Machines tab.
2. Click the Actions drop-down menu at the top of the Machines tab and then click **Remove Machines**.
3. Select one of the options described in the following table.

Option	Description
Keep Recovery Points	To keep all currently stored recovery points for this cluster.
Remove Recovery Points	To remove all currently stored recovery points for this cluster from the repository.

Viewing a Cluster or Node Report

You can create and view compliance and errors reports about AppAssure 5 activities for your cluster and individual nodes. The reports include AppAssure 5 activity information about the cluster, node, and shared volumes.

For more information about AppAssure 5 reporting, see Chapter 6, "Reporting". For more information about the exporting and printing options located in the reports toolbar, see "About the Reports Toolbar" on page 208.

To view a cluster or node report

1. Do one of the following:
   - In the Core Console, click the Machines tab, and then select the cluster or node for which you wish to create a report.
   - Or, in the left navigation area, select the cluster or node for which you wish to create a report.
2. Click the Tools tab and, under the Reports menu, select one of the following options:
   - Compliance Report
   - Errors Report

3. 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 AppAssure 5 Core or Agent was deployed.

4. In the End Time drop-down calendar, select an end date, and then enter an end time for the report.

5. Click **Generate Report**. The report results appear in the page.

   If the report spans multiple pages, you can click the page numbers or the arrow buttons at the top of the report results to page through the results.

6. To export the report results to one of the available formats–PDF, XLS, XLSX, RTF, MHT, HTML, TXT, CSV, or image–select the format for export from the drop-down list, and then do one of the following:

   - Click the first Save icon to export a report and save it to the disk.
   - Click the second Save icon to export a report and show it in a new Web browser window.

7. To print the report results, do one of the following:

   - Click the first Printer icon to print the entire report.
   - Click the second Printer icon to print the current page of the report.

6 Reporting

This chapter provides an overview of reporting available in AppAssure 5. It consists of the following topics:

About Reports

AppAssure 5 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 email the URL for others to view the report with a Web browser.

For information about generating a report, see "Generating a Report for a Core or Agent" on page 210. For information about the generating a report for multiple cores in the Central Management Console, see "Generating a Report from the Central Management Console" on page 211. For information about generating cluster reports, see "Viewing a Cluster or Node Report" on page 205.

About Compliance Reports

Compliance Reports are available for the AppAssure 5 Core and AppAssure 5 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 appears 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

For information about how to generate a report, see "Generating a Report for a Core or Agent" on page 210.

About Errors Reports

Errors Reports are subsets of the Compliance Reports and are available for AppAssure 5 Cores and AppAssure 5 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

For information about how to generate a report, see "Generating a Report for a Core or Agent" on page 210.

About the Core Summary Report

The Core Summary Report includes information about the repositories on the selected AppAssure 5 Core and about the agents protected by that core. The information appears as two summaries within one report.

For information on how to generate a Core Summary Report, see "Generating a Report for a Core or Agent" on page 210.

Repositories Summary

The Repositories portion of the Core Summary Report includes data for the repositories located on the selected AppAssure 5 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 AppAssure 5 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

Complete the steps in the following procedure to generate a report for an AppAssure 5 Core or AppAssure 5 Agent.

To generate a report for a core or agent

1. Navigate to the AppAssure 5 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 in the following table.

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. For more information about the toolbar, see "About the Reports Toolbar" on page 208.

About the Central Management Console Core Reports

AppAssure 5 lets you generate and view compliance, error, and summary information for multiple AppAssure 5 Cores. Details about the Cores are presented in column views with the same categories described in the sections "About Compliance Reports" on page 208, "About Errors Reports" on page 209, and "About the Core Summary Report" on page 209.

For information on how to generate a report for multiple cores, see "Generating a Report from the Central Management Console" on page 211.

Generating a Report from the Central Management Console

Complete the following procedure to generate a report for multiple AppAssure 5 Cores from the Central Management Console.

To generate a report from the Central Management Console

1. From the Central Management Console Welcome screen, click 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
   - Errors Report
3. From the left navigation area, select the AppAssure 5 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 were 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. For more information about the toolbar, see "About the Reports Toolbar" on page 208.

A Scripting

This appendix describes the scripts that can be used by administrators at designated occurrences in AppAssure 5 for Windows and Linux. It includes the following topics:

Scripting in AppAssure 5

AppAssure 5 enables administrators to automate the administration and management of resources at certain occurrences through the execution of commands and scripts. AppAssure 5 supports the use of PowerShell scripting for Windows and Bourne Shell scripting for Linux. For more information on how using PowerShell scripts see "About PowerShell Scripting in AppAssure 5" on page 214, "Sample PowerShell Scripts" on page 226, "Input Parameters for PowerShell Scripting" on page 215, and "Sample Bourne Shell Scripts" on page 240

About PowerShell Scripting in AppAssure 5

Windows PowerShell is a Microsoft .NET Framework-connected environment designed for administrative automation. AppAssure 5 includes comprehensive client software development kits (SDKs) for PowerShell scripting that lets administrative users execute user-provided PowerShell scripts at designated occurrences; for example, before or after a snapshot, attachability and mountability checks, and so on. Administrators can execute scripts from both the AppAssure 5 Core and the Agent. Scripts can accept parameters, and the output of a script is written to core and agent log files.


Note

For nightly jobs, you should preserve one script file and the JobType input parameter to distinguish between nightly jobs.

Script files are located in the %ALLUSERSPROFILE%\AppRecovery\Scripts folder.

- In Windows 7, the path to locate the %ALLUSERSPROFILE% folder is: C:\ProgramData.
- In Windows 2003, the path to locate the folder is: Documents and Settings\All Users\Application Data\.


Note

Windows PowerShell is required and must be installed and configured prior to using and executing AppAssure 5 scripts.

Prerequisites for PowerShell Scripting

Before using and executing the PowerShell scripts for AppAssure 5, you must have Windows PowerShell 3.0 installed.


Note

Make sure to place the powershell.exe.config file in the PowerShell home directory. For example, C:\WindowsPowerShell\powershell.exe.

powershell.exe.config

```
<?xml version="1.0"?>

<configuration>

     <startup useLegacyV2RuntimeActivationPolicy="true">
        <supportedRuntime version="v4.0.30319"/>
        <supportedRuntime version="v2.0.50727"/>
     </startup>


</configuration>
```

Testing PowerShell Scripts

If you want to test the scripts you plan to run, you can do so by using the PowerShell graphical editor, powershell_is. You also need to add the configuration file, powershell_ise.exe.config to the same folder the configuration file, powershell.exe.config.

The configuration file, powershell_ise.exe.config must have the same content as that of the powershell.exe.config file.

If the pre-PowerShell or post-PowerShell script fails, the job also fails.

Input Parameters for PowerShell Scripting

All available input parameters are used in sample scripts. The parameters are described in the following tables.

Script files must possess the same name as the sample script files.

AgentProtectionStorageConfiguration (namespace Replay.Common.Contracts.Agents)

The following table presents the available objects for the AgentProtectionStorageConfiguration parameter.

Method	Description
public Guid RepositoryId { get; set; }	Gets or sets the ID of the repository where the agent recovery points are stored.
public string EncryptionKeyId { get; set; }	Gets or sets the ID of the encryption key for this agent's recovery points. An empty string means no encryption.

AgentTransferConfiguration (namespace Replay.Common.Contracts.Transfer)

The following table presents the available objects for the AgentTransferConfiguration parameter.

Method	Description
public uint MaxConcurrentStreams { get; set; }	Gets or sets the maximum number of concurrent TCP connections the Core establishes to the agent for transferring data.
public uint MaxTransferQueueDepth { get; set; }	When a range of blocks are read from a transfer stream, that range is placed on a producer or consumer queue, where a consumer thread reads it and writes it to the epoch object. If the repository writes slower than the network reads, this queue fills up. The point at which the queue is full and reads stop, is the max transfer queue depth.
public uint MaxConcurrentWrites { get; set; }	Gets or sets the maximum number of block write operations to have outstanding on an epoch at any given time. If additional blocks are received when this many block writes are outstanding, those additional blocks are ignored until one of the outstanding writes finishes.
public ulong MaxSegmentSize { get; set; }	Gets or sets the maximum number of contiguous blocks to transfer in a single request. Depending on testing, higher or lower values may be optimal.
public Priority Priority { get; set; }	Gets or sets the priority for transfer request.
public int MaxRetries { get; set; }.	Gets or sets the maximum number of times a failed transfer should be retried before it is presumed failed.
public Guid ProviderId{ get; set; }	Gets or sets the GUID of the VSS provider to use for snapshots on this host. Administrators typically accept the default.
public Collection<ExcludedWriter> ExcludedWriterIds { get; set; }	Gets or sets the collection of VSS writer IDs, which should be excluded from this snapshot. The writer ID is determined by the name of the writer. This name is for documentation purposes only and does not have to exactly match the name of the writer.
public ushort TransferDataServerPort { get; set; }	Gets or sets a value containing the TCP port upon which to accept connections from the Core for the actual transfer of data from the Agent to the Core. The Agent attempts to listen on this port, but if the port is in use, the Agent can use a different port instead. The Core should use the port number specified in the BlockHashesUri and BlockDataUri properties of the VolumeSnapshotInfo object for each snapped volume.
public TimeSpan SnapshotTimeout { get; set; }	Gets or sets the amount of time to wait for a VSS snapshot operation to complete before giving up and timing out.

Method	Description
public TimeSpan TransferTimeout { get; set; }	Gets or sets the amount of time to wait for further contact from the Core before abandoning the snapshot.
public TimeSpan NetworkReadTimeout { get; set; }	Gets or sets the timeout for network read operations related to this transfer.
public TimeSpan NetworkWriteTimeout { get; set; }	Gets or sets the timeout for network write operations related to this transfer.

BackgroundJobRequest (namespace Replay.Core.Contracts.BackgroundJobs)

The following table presents the available objects for the BackgroundJobRequest parameter.

Method	Description
public Guid AgentId { get; set; }	Gets or sets the ID of the Agent.
public bool IsNightlyJob { get; set; }	Gets or sets the value indicating whether the background job is a nightly job.
public virtual bool InvolvesAgentId(Guid agentId)	Determines the value indicating whether the concrete agent is involved in job.

ChecksumCheckJobRequest (namespace Replay.Core.Contracts.Exchange.ChecksumChecks)

Inherits its values from the parameter, DatabaseCheckJobRequestBase.

DatabaseCheckJobRequestBase (namespace Replay.Core.Contracts.Exchange)

Inherits its values from the parameter, BackgroundJobRequest.

ExportJobRequest (namespace Replay.Core.Contracts.Export)

Inherits its values from the parameter, BackgroundJobRequest.
The following table presents the available objects for the ExportJobRequest parameter.

Method	Description
public uint RamInMegabytes { get; set; }	Gets or sets the memory size for the exported VM. Set to zero (0) to use the memory size of the source machine.
public VirtualMachineLocation Location { get; set; }	Gets or sets the target location for this export. This is an abstract base class.
public VolumeImageIdsCollection VolumeImageIds { get; private set; }	Gets or sets the volume images to include in the VM export.
public ExportJobPriority Priority { get; set; }	Gets or sets the priority for export request.

NightlyAttachabilityJobRequest (namespace Replay.Core.Contracts.Sql)

Inherits its values from the parameter, BackgroundJobRequest.

RollupJobRequest (namespace Replay.Core.Contracts.Rollup)

Inherits its values from the parameter, BackgroundJobRequest.

TakeSnapshotResponse (namespace Replay.Agent.Contracts.Transfer)

The following table presents the available objects for the TakeSnapshotResponse parameter.

Method	Description
public Guid SnapshotSetId { get; set; }	Gets or sets the GUID assigned by VSS to this snapshot.
public VolumeSnapshotInfoDictionary VolumeSnapshots { get; set; }	Gets or sets the collection of snapshot info for each volume included in the snap.

TransferJobRequest (namespace Replay.Core.Contracts.Transfer)

Inherits its values from the parameter, BackgroundJobRequest.
The following table presents the available objects for the TransferJobRequest parameter.

Method	Description
public VolumeNameCollection VolumeNames { get; set; }	Gets or sets the collection of names for transfer. VolumeNames is a data structure that contains the following data: ■ GuidName - The Guid associated with the volume, used as the name if a DisplayName is not set. ■ DisplayName - The displayed name of the volume.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer. The available values are: ■ Unknown ■ Copy ■ Full
public AgentTransferConfiguration TransferConfiguration { get; set; }	Gets or sets the transfer configuration.
public AgentProtectionStorageConfiguration StorageConfiguration { get; set; }	Gets or sets the storage configuration.
public string Key { get; set; }	Generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.

Method	Description
public bool ForceBaseImage { get; set; }	Gets or sets the value indicating whether the base image was forced or not.
public bool IsLogTruncation { get; set; }	Gets or sets the value indicating whether the job is log truncation or not.

TransferPrescriptParameter (namespace Replay.Common.Contracts.PowerShellExecution)

The following table presents the available objects for the TransferPrescript parameter.

Method	Description
public VolumeNameCollection VolumeNames (get; set;)	Gets or sets the collection of volume names for transfer. VolumeNames is a data structure that contains the following data: ■ GuidName - The Guid associated with the volume, used as the name if a DisplayName is not set. ■ DisplayName - The displayed name of the volume.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer.ShadowCopyType is an enumeration with values. The available values are: ■ Unknown ■ Copy ■ Full

Method	Description
public AgentTransferConfiguration TransferConfiguration { get; set; }	Gets or sets the transfer configuration. AgentTransferConfiguration is an object which will have the following data: ■ MaxConcurrentStreams - the maximum number of concurrent TCP connections the core will establish to the agent for transferring data ■ MaxTransferQueueDepth - the maximum number of block extents which can be queued up for writing ■ MaxConcurrentWrites - the maximum number of block write operations to have outstanding on an epoch at any given time. If additional blocks are received when this many block writes are outstanding, those additional blocks will be ignored until one of the outstanding blocks gets written. ■ MaxSegmentSize - the maximum number of contiguous blocks to transfer in a single request ■ Priority - An object which will have the following data: □ Undefined □ One □ Two □ Three □ Four □ Five □ Six □ Seven □ Eight □ Nine □ Ten □ Highest (which is equal to One) □ Lowest (which is equal to Ten) □ Default (which is equal to Five) ■ MaxRetries - the maximum number of times a failed transfer should be retried before it is presumed failed ■ UseDefaultMaxRetries - a value indicating that the maximum number of retries is the default value ■ ProviderId - the GUID of the VSS provider to use for snapshots on this host. Pretty much everyone will want to accept the default.

Method	Description
public AgentTransferConfiguration TransferConfiguration { get; set; } (cont.)	■ ExcludedWriterIds - collection of VSS writer IDs which should be excluded from this snapshot. The writer ID is keyed by the name of the writer. This name is for documentation purposes only and does not have to exactly match the actual name of the writer. ■ TransferDataServerPort - a value containing the TCP port upon which to accept connections from the core for the actual transfer of data from the agent to the core. ■ SnapshotTimeout - the amount of time to wait for a VSS snapshot operation to complete before giving up and timing out. ■ TransferTimeout - the amount of time to wait for further contact from the core before abandoning the snapshot. ■ NetworkReadTimeout - the timeout for network read operations related to this transfer. ■ NetworkWriteTimeout - the timeout for network write operations related to this transfer. ■ InitialQueueSize - a size of initial queue of requests. ■ MinVolumeFreeSpacePercents - a minimal amount of free space on a volume in percent. ■ MaxChangeLogsSizePercents - a maximum size of driver change logs as part of volume capacity measured in percent. ■ EnableVerification - a value indicating whether diagnostic verification of each block sent to Core should be performed.
public string Key { get; set; }	The Key method generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.
public bool ForceBaseImage { get; set; }	Gets or sets the value indicating whether the transfer was a forced base image capture.
public bool IsLogTruncation { get; set; }	Gets or sets the value indicating whether logging is being truncated.
public uint LatestEpochSeenByCore { get; set; }	Gets or sets latest epoch value. The LatestEpochSeenByCore method is the ordinal number of the most recent snapshot taken by the Core. This is the 'epoch number' assigned by the filter driver to this particular snapshot at the moment it was taken with VSS.

TransferPostscriptParameter (namespace Replay.Common.Contracts.PowerShellExecution)

The following table presents the available objects for the TransferPostscript parameter.

Method	Description
public VolumeNameCollection VolumeNames (get; set;)	Gets or sets the collection of volume names for transfer. VolumeNames is a data structure that contains the following data: ■ GuidName - The Guid associated with the volume, used as the name if a DisplayName is not set. ■ DisplayName - The displayed name of the volume.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer.ShadowCopyType is an enumeration with values. The available values are: ■ Unknown ■ Copy ■ Full

Method	Description
public AgentTransferConfiguration TransferConfiguration { get; set; }	Gets or sets the transfer configuration. AgentTransferConfiguration is an object which will have the following data: ■ MaxConcurrentStreams - the maximum number of concurrent TCP connections the core will establish to the agent for transferring data ■ MaxTransferQueueDepth - the maximum number of block extents which can be queued up for writing ■ MaxConcurrentWrites - the maximum number of block write operations to have outstanding on an epoch at any given time. If additional blocks are received when this many block writes are outstanding, those additional blocks will be ignored until one of the outstanding blocks gets written. ■ MaxSegmentSize - the maximum number of contiguous blocks to transfer in a single request ■ Priority - An object which will have the following data: □ "Undefined □ "One □ "Two □ "Three □ "Four □ "Five □ "Six □ "Seven □ "Eight □ "Nine □ "Ten □ "Highest (which is equal to One) □ "Lowest (which is equal to Ten) □ "Default (which is equal to Five) ■ MaxRetries - the maximum number of times a failed transfer should be retried before it is presumed failed ■ UseDefaultMaxRetries - a value indicating that the maximum number of retries is the default value ■ ProviderId - the GUID of the VSS provider to use for snapshots on this host. Pretty much everyone will want to accept the default

Method	Description
public AgentTransferConfiguration TransferConfiguration { get; set; } (cont.)	■ ExcludedWriterIds - collection of VSS writer IDs which should be excluded from this snapshot. The writer ID is keyed by the name of the writer. This name is for documentation purposes only and does not have to exactly match the actual name of the writer. ■ TransferDataServerPort - a value containing the TCP port upon which to accept connections from the core for the actual transfer of data from the agent to the core. ■ SnapshotTimeout - the amount of time to wait for a VSS snapshot operation to complete before giving up and timing out. ■ TransferTimeout - the amount of time to wait for further contact from the core before abandoning the snapshot. ■ NetworkReadTimeout - the timeout for network read operations related to this transfer. ■ NetworkWriteTimeout - the timeout for network write operations related to this transfer. ■ InitialQueueSize - a size of initial queue of requests. ■ MinVolumeFreeSpacePercents - a minimal amount of free space on a volume in percent. ■ MaxChangeLogsSizePercents - a maximum size of driver change logs as part of volume capacity measured in percent. ■ EnableVerification - a value indicating whether diagnostic verification of each block sent to Core should be performed.
public AgentProtectionStorageConfiguration StorageConfiguration { get; set; }	Gets or sets the storage configuration The AgentProtectionStorageConfiguration object contains the following data: ■ RepositoryId - the name of the repository where this agent's recovery points will be stored ■ EncryptionKeyId - the ID of the encryption key for this agent's recovery points. An empty string means no encryption
public string Key { get; set; }	The Key method generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.
public bool ForceBaseImage { get; set; }	Gets or sets the value indicating whether the transfer was a forced base image capture.
public bool IsLogTruncation { get; set; }	Gets or sets the value indicating whether logging is being truncated.

Method	Description
public uint LatestEpochSeenByCore { get; set; }	Gets or sets latest epoch value. The LatestEpochSeenByCore method is the ordinal number of the most recent snapshot taken by the Core. This is the 'epoch number' assigned by the filter driver to this particular snapshot at the moment it was taken with VSS.
public Guid SnapshotSetId { get; set; }	Gets or sets the GUID assigned by VSS to this snapshot.
public VolumeSnapshotInfoDictionary VolumeSnapshots { get; set; }	Gets or sets the collection of snapshot info for each volume included in the snap.

VirtualMachineLocation (namespace Replay.Common.Contracts.Virtualization)

The following table presents the available objects for the VirtualMachineLocation parameter.

Method	Description
public string Description { get; set;}	Gets or sets a human-readable description of this location.
public string Method { get; set;}	Gets or sets the name of the VM.

VolumeImageIdsCollection (namespace Replay.Core.Contracts.RecoveryPoints)

Inherits its values from the parameter, System.Collections.ObjectModel.Collection<string>.

VolumeName (namespace Replay.Common.Contracts.Metadata.Storage)

The following table presents the available objects for the VolumeName parameter.

Method	Description
public string GuidName { get; set;}	Gets or sets the ID of the volume.
public string DisplayName { get; set;}	Gets or sets the name of the volume.
public string UrlEncode()	Gets a URL-encoded version of the name which can be passed cleanly on a URL. **NOTE:** A known issue exists in .NET 4.0 WCF (https://connect.microsoft.com/VisualStudio/feedback/ViewFeedback.aspx?FeedbackID=413312), which prevents path escape characters from working correctly in a URI template. Because a volume name contains both '\' and '?', you must replace the special characters '\' and '?' with other special characters.
public string GetMountName()	Returns a name for this volume that is valid for mounting volume image to some folder.

VolumeNameCollection (namespace Replay.Common.Contracts.Metadata.Storage)

Inherits its values from the parameter, System.Collections.ObjectModel.Collection<VolumeName>.
The following table presents the available objects for the VolumeNameCollection parameter.

Method	Description
public override bool Equals(object obj)	Determines whether this instance and a specified object, which must also be a VolumeNameCollection object, have the same value. (Overrides Object.Equals(Object).)
public override int GetHashCode()	Returns the hash code for this VolumeNameCollection. (Overrides Object.GetHashCode().)

VolumeSnapshotInfo (namesapce Replay.Common.Contracts.Transfer)

The following table presents the available objects for the VolumeSnapshotInfo parameter.

Method	Description
public Uri BlockHashesUri { get; set;}	Gets or sets the URI at which the MD5 hashes of volume blocks can be read.
public Uri BlockDataUri { get; set;}	Gets or sets the URI at which the volume data blocks can be read.

VolumeSnapshotInfoDictionary (namespace Replay.Common.Contracts.Transfer)

Inherits its values from the parameter, System.Collections.Generic.Dictionary<VolumeName, VolumeSnapshotInfo>.

Sample PowerShell Scripts

The following sample scripts are provided to assist administrative users in executing PowerShell scripts. The sample scripts include:

- PreTransferScript.ps1
- PostTransferScript.ps1
- PreExportScript.ps1
- PostExportScript.ps1
- PreNightlyJobScript.ps1
- PostNightlyJobScript.ps1

PreTransferScript.ps1

The PreTransferScript is executed on the agent side prior to transferring a snapshot.

Sample PreTransferScript

```
# receiving parameter from transfer job

param([object]$TransferPrescriptParameter)


# building path to Agent's Common.Contracts.dll and loading this assembly

$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Agent 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null


# Converting input parameter into specific object

$TransferPrescriptParameterObject = $TransferPrescriptParameter -as
[Replay.Common.Contracts.PowerShellExecution.TransferPrescriptParameter];


# Working with input object. All echo's are logged

if($TransferPrescriptParameterObject -eq $null) {

        echo 'TransferPrescriptParameterObject parameter is null'

}

else {

        echo
'TransferConfiguration:'$TransferPrescriptParameterObject.TransferConfigu
ration

        echo 'StorageConfiguration:'
$TransferPrescriptParameterObject.StorageConfiguration

}
```

PostTransferScript.ps1

The PostTransferScript is executed on the agent side after transferring a snapshot.

Sample PostTransferScript

```
# receiving parameter from transfer job

param([object] $TransferPostscriptParameter)


# building path to Agent's Common.Contracts.dll and loading this assembly


$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Agent 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null


# Converting input parameter into specific object

$TransferPostscriptParameterObject = $TransferPostscriptParameter -as
[Replay.Common.Contracts.PowerShellExecution.TransferPostscriptParameter]
;


# Working with input object. All echo's are logged

if($TransferPostscriptParameterObject -eq $null) {

        echo 'TransferPostscriptParameterObject parameter is null'

}

else {

echo 'VolumeNames:' $TransferPostscriptParameterObject.VolumeNames

        echo 'ShadowCopyType:'
$TransferPostscriptParameterObject.ShadowCopyType

    echo 'ForceBaseImage:'
$TransferPostscriptParameterObject.ForceBaseImage

    echo 'IsLogTruncation:'
$TransferPostscriptParameterObject.IsLogTruncation

}
```

PreExportScript.ps1

The PreExportScript is executed on the Core side prior to any export job.

Sample PreExportScript

```
# receiving parameter from export job

param([object]$ExportJobRequest)


# building path to Core's Common.Contracts.dll and loading this assembly

$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Core 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'CoreService\Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null


# Converting input parameter into specific object

$ExportJobRequestObject = $ExportJobRequest -as
[Replay.Core.Contracts.Export.ExportJobRequest]


# Working with input object. All echo's are logged

if($ExportJobRequestObject -eq $null) {

        echo 'ExportJobRequestObject parameter is null'

}

else {

        echo 'Location:' $ExportJobRequestObject.Location

        echo 'Priority:' $ExportJobRequestObject.StorageConfiguration

}
```

PostExportScript.ps1

The PostExportScript is executed on the Core side after any export job.

NOTE: There are no input parameters for the PostExportScript when used to execute once on the exported agent after initial startup. The regular agent should contain this script in the PowerShell script folder as PostExportScript.ps1.

Sample PostExportScript

```
# receiving parameter from export job

param([object]$ExportJobRequest)


# building path to Core's Common.Contracts.dll and loading this assembly

$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Core 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'CoreService\Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null

$regVal2 = $regLM.GetValue('InstallLocation')

$regVal2 = $regVal2 + 'CoreService\Common.Contracts.dll'


# Converting input parameter into specific object

$ExportJobRequestObject = $ExportJobRequest -as
[Replay.Core.Contracts.Export.ExportJobRequest]


# Working with input object. All echo's are logged

if($ExportJobRequestObject -eq $null) {

        echo 'ExportJobRequestObject parameter is null'

}

else {

        echo 'VolumeImageIds:' $ExportJobRequestObject.VolumeImageIds

        echo 'RamInMegabytes:' $ExportJobRequestObject.RamInMegabytes

}
```

PreNightlyJobScript.ps1

The PreNightlyJobScript is executed before every nighty job on Core side. It has $JobClassName parameter, that helps to handle those child jobs separately.

Sample PreNightlyJobScript

```
# receiving parameters from Nightlyjob

param([System.String]$JobClassMethod ,
[object]$NightlyAttachabilityJobRequest, [object]$RollupJobRequest,
[object]$Agents, [object]$ChecksumCheckJobRequest,
[object]$TransferJobRequest, [int]$LatestEpochSeenByCore)



# building path to Core's Common.Contracts.dll and loading this assembly

$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Core 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'CoreService\Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null



# Nightlyjob has four child jobs: NightlyAttachability Job, Rollup Job,
Checksum Check Job and Log Truncation Job. All of them are triggering the
script, and $JobClassMethod (contain job name that calls the script) helps
to handle those child jobs separately

switch ($JobClassMethod) {



# working with NightlyAttachability Job

    NightlyAttachabilityJob {

        $NightlyAttachabilityJobRequestObject =
$NightlyAttachabilityJobRequest -as
[Replay.Core.Contracts.Sql.NightlyAttachabilityJobRequest];

        echo 'Nightly Attachability job results:';

        if($NightlyAttachabilityJobRequestObject -eq $null) {

            echo 'NightlyAttachabilityJobRequestObject parameter is null';

        }

        else {
```

```
                echo 'AgentId:'
$NightlyAttachabilityJobRequestObject.AgentId;

                echo 'IsNightlyJob:'
$NightlyAttachabilityJobRequestObject.IsNightlyJob;

        }

        break;

    }


# working with Rollup Job

    RollupJob {

        $RollupJobRequestObject = $RollupJobRequest -as
[Replay.Core.Contracts.Rollup.RollupJobRequest];

        echo 'Rollup job results:';

        if($RollupJobRequestObject -eq $null) {

                echo 'RollupJobRequestObject parameter is null';

        }

        else {

                echo 'SimultaneousJobsCount:'
$RollupJobRequestObject.SimultaneousJobsCount;

                echo 'AgentId:' $RollupJobRequestObject.AgentId;

               echo 'IsNightlyJob:' $RollupJobRequestObject.IsNightlyJob;

        }

        $AgentsCollection = $Agents -as
"System.Collections.Generic.List``1[System.Guid]"

        if($AgentsCollection -eq $null) {

            echo 'AgentsCollection parameter is null';

        }

        else {

            echo 'Agents GUIDs:'

            foreach ($a in $AgentsCollection) {

                echo $a

            }

        }

        break;

    }
```

```
# working with Checksum Check Job

    ChecksumCheckJob {

        $ChecksumCheckJobRequestObject = $ChecksumCheckJobRequest -as
[Replay.Core.Contracts.Exchange.ChecksumChecks.ChecksumCheckJobRequest];

        echo 'Exchange checksumcheck job results:';

        if($ChecksumCheckJobRequestObject -eq $null) {

                echo 'ChecksumCheckJobRequestObject parameter is null';

        }

        else {

                echo 'RecoveryPointId:'
$ChecksumCheckJobRequestObject.RecoveryPointId;

                echo 'AgentId:' $ChecksumCheckJobRequestObject.AgentId;

                echo 'IsNightlyJob:'
$ChecksumCheckJobRequestObject.IsNightlyJob;

        }

        break;

    }


# working with Log Truncation Job

    TransferJob {

        $TransferJobRequestObject = $TransferJobRequest -as
[Replay.Core.Contracts.Transfer.TransferJobRequest];

        echo 'Transfer job results:';

        if($TransferJobRequestObject -eq $null) {

                echo 'TransferJobRequestObject parameter is null';

        }

        else {

                echo 'TransferConfiguration:'
$TransferJobRequestObject.TransferConfiguration;

                echo 'StorageConfiguration:'
$TransferJobRequestObject.StorageConfiguration;

        }

        echo 'LatestEpochSeenByCore:' $LatestEpochSeenByCore;

        break;
```

```
    }
}
```

PostNightlyJobScript.ps1

The PostNightlyJobScript is executed after every nighty job on Core side. It has $JobClassName parameter, that helps to handle those child jobs separately.

Sample PostNightlyJobScript

```
# receiving parameters from Nightlyjob

param([System.String]$JobClassMethod ,
[object]$NightlyAttachabilityJobRequest, [object]$RollupJobRequest,
[object]$Agents, [object]$ChecksumCheckJobRequest,
[object]$TransferJobRequest, [int]$LatestEpochSeenByCore,
[object]$TakeSnapshotResponse)


# building path to Core's Common.Contracts.dll and loading this assembly

$regLM = [Microsoft.Win32.Registry]::LocalMachine

$regLM =
$regLM.OpenSubKey('SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\Ap
pRecovery Core 5')

$regVal = $regLM.GetValue('InstallLocation')

$regVal = $regVal + 'CoreService\Common.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal)  |  out-null

$regVal2 = $regLM.GetValue('InstallLocation')

$regVal2= $regVal2 + 'CoreService\Core.Contracts.dll'

[System.Reflection.Assembly]::LoadFrom($regVal2) | out-null


# Nightlyjob has four child jobs: NightlyAttachability Job, Rollup Job,
Checksum Check Job and Log Truncation Job. All of them are triggering the
script, and $JobClassMethod (contain job name that calls the script) helps
to handle those child jobs separately

switch ($JobClassMethod) {


# working with NightlyAttachability Job

    NightlyAttachabilityJob {
```

```
        $NightlyAttachabilityJobRequestObject =
$NightlyAttachabilityJobRequest -as
[Replay.Core.Contracts.Sql.NightlyAttachabilityJobRequest];

        echo 'Nightly Attachability job results:';

        if($NightlyAttachabilityJobRequestObject -eq $null) {

            echo 'NightlyAttachabilityJobRequestObject parameter is null';

        }

        else {

                echo 'AgentId:'
$NightlyAttachabilityJobRequestObject.AgentId;

                echo 'IsNightlyJob:'
$NightlyAttachabilityJobRequestObject.IsNightlyJob;

        }

        break;

    }


# working with Rollup Job

    RollupJob {

        $RollupJobRequestObject = $RollupJobRequest -as
[Replay.Core.Contracts.Rollup.RollupJobRequest];

        echo 'Rollup job results:';

        if($RollupJobRequestObject -eq $null) {

                echo 'RollupJobRequestObject parameter is null';

        }

        else {

                echo 'SimultaneousJobsCount:'
$RollupJobRequestObject.SimultaneousJobsCount;

                echo 'AgentId:' $RollupJobRequestObject.AgentId;

               echo 'IsNightlyJob:' $RollupJobRequestObject.IsNightlyJob;

        }

        $AgentsCollection = $Agents -as
"System.Collections.Generic.List``1[System.Guid]"

        if($AgentsCollection -eq $null) {

            echo 'AgentsCollection parameter is null';

        }

        else {
```

```
            echo 'Agents GUIDs:'
            foreach ($a in $AgentsCollection) {
                echo $a
            }
        }
        break;
    }

# working with Checksum Check Job
    ChecksumCheckJob {
        $ChecksumCheckJobRequestObject = $ChecksumCheckJobRequest -as
[Replay.Core.Contracts.Exchange.ChecksumChecks.ChecksumCheckJobRequest];
        echo 'Exchange checksumcheck job results:';
        if($ChecksumCheckJobRequestObject -eq $null) {
                echo 'ChecksumCheckJobRequestObject parameter is null';
        }
        else {
                echo 'RecoveryPointId:'
$ChecksumCheckJobRequestObject.RecoveryPointId;
                echo 'AgentId:' $ChecksumCheckJobRequestObject.AgentId;
            echo 'IsNightlyJob:'
$ChecksumCheckJobRequestObject.IsNightlyJob;
        }
        break;
    }

# working with Log Truncation Job
    TransferJob {
        $TransferJobRequestObject = $TransferJobRequest -as
[Replay.Core.Contracts.Transfer.TransferJobRequest];
        echo 'Transfer job results:';
        if($TransferJobRequestObject -eq $null) {
                echo 'TransferJobRequestObject parameter is null';
        }
```

```
        else {

                echo 'TransferConfiguration:'
$TransferJobRequestObject.TransferConfiguration;

                echo 'StorageConfiguration:'
$TransferJobRequestObject.StorageConfiguration;

        }

        echo 'LatestEpochSeenByCore:' $LatestEpochSeenByCore;

        $TakeSnapshotResponseObject = $TakeSnapshotResponse -as
[Replay.Agent.Contracts.Transfer.TakeSnapshotResponse];

        if($TakeSnapshotResponseObject -eq $null) {

                echo 'TakeSnapshotResponseObject parameter is null';

        }

        else {

                echo 'ID of this transfer session:'
$TakeSnapshotResponseObject.Id;

                echo 'Volumes:' $TakeSnapshotResponseObject.Volumes;

        }

        break;

    }

}
```

About Bourne Shell Scripting in AppAssure 5

Bourne shell (sh) or Bourne Again Shell (BASH) is a shell language or command-line interpreter for Unix-based operating systems and is used in AppAssure 5 with Linux to customize environments and specify certain operations to occur in a predetermined sequence. The .sh is the file extension and naming convention for Bourne shell files.

Using the pre and post transfer script hooks, you can perform system operations before and after a transfer. For example, you may want to disable a certain cronjob while a transfer is occurring and enable it once the transfer has finished. Another example could include the need to execute commands to flush application specific data to disk. The contents is written to a temporary file and executed using exec. This will cause the script to be executed using the interpreter defined in the first line of the script, for example, (#!/usr/bin/env bash) or the default shell as defined by the $SHELL environment variable if that isn't present. Depending on your preference, you can substitute and use any interpreter, for example, zsh, tcsh, and so on in the #! line of the script to use whatever their preference is, should it vary from the default shell.

You can add available objects from the TransferPrescript parameter or add your own commands to the PreTransferScript.sh and PostTransfer.sh scripts to customize them.

Prerequisites for Bourne Shell Scripting

All scripts must be named PreTransferScript.sh and PostTransfer.sh and have to reside in the /opt/appassure/scripts/ directory.

Testing Bourne Shell Scripting

You can test the scripts you want to run by using the editor for the script (.sh) files.


Note

If the pre-Bourne Shell or post-Bourne Shell scripts fail, the job will also fails. Information about the job is available in the /var/log/appassure/appassure.log file.

Input Parameters for Bourne Shell Scripting

The parameters for Bourne Shell scripting in AppAssure 5 are described in the following tables.

TransferPrescriptParameter

The following table presents the available objects for the TransferPrescript parameter.

Method	Description
public VolumeNameCollection VolumeNames (get; set;)	Gets or sets the collection of volume names for transfer. VolumeNames is a data structure that contains the following data: ■ GuidName - The Guid associated with the volume, used as the name if a DisplayName is not set. ■ DisplayName - The displayed name of the volume.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer. ShadowCopyType is an enumeration with values. The available values are: ■ Unknown ■ Copy ■ Full

Method	Description
public string Key { get; set; }	The Key method generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.
public bool ForceBaseImage { get; set; }	Gets or sets the value indicating whether the transfer was a forced base image capture.
public bool IsLogTruncation { get; set; }	Gets or sets the value indicating whether logging is being truncated.
public uint LatestEpochSeenByCore { get; set; }	Gets or sets latest epoch value. The LatestEpochSeenByCore method is the ordinal number of the most recent snapshot taken by the Core. This is the 'epoch number' assigned by the filter driver to this particular snapshot at the moment it was taken with VSS.

TransferPostscriptParameter

The following table presents the available objects for the TransferPostscript parameter.

Method	Description
public VolumeNameCollection VolumeNames (get; set;)	Gets or sets the collection of volume names for transfer. VolumeNames is a data structure that contains the following data: ■ GuidName - The Guid associated with the volume, used as the name if a DisplayName is not set. ■ DisplayName - The displayed name of the volume.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer.ShadowCopyType is an enumeration with values. The available values are: ■ Unknown ■ Copy ■ Full
public string Key { get; set; }	The Key method generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.
public bool ForceBaseImage { get; set; }	Gets or sets the value indicating whether the transfer was a forced base image capture.

Method	Description
public bool IsLogTruncation { get; set; }	Gets or sets the value indicating whether logging is being truncated.
public uint LatestEpochSeenByCore { get; set; }	Gets or sets latest epoch value. The LatestEpochSeenByCore method is the ordinal number of the most recent snapshot taken by the Core. This is the 'epoch number' assigned by the filter driver to this particular snapshot at the moment it was taken with VSS.

Sample Bourne Shell Scripts

The following sample scripts are provided to assist administrative users in executing Bourne Shell scripts for agents. You can use the sample scripts and customize them as needed. The sample scripts for agents include:

- PreTransferScript.sh
- PostTransferScript.sh

The agent uses the 'exec' shell command to launch the scrip. You can indicate which interpreter should run the script by defining that information in the first line of the script, which is similar to what you define for any normal script being executed from the command line. If you don't specify the interpreter, the default shell will interpret the script. If you choose something other than the default shell, you will have to ensure that the specified interpreter is available on all protected machines.

PreTransferScript.sh

The PreTransferScript is executed on the agent side prior to transferring a snapshot.

The following script stores the values from input parameters in the Pre(Post)TransferScriptResult.txt which is located and stored in the root home directory.

Sample PreTransferScript

```
#!/bin/bash

echo
"TransferPrescriptParameter_VolumeNames=$TransferPrescriptParameter_Volum
eNames

TransferPrescriptParameter_ShadowCopyType=$TransferPrescriptParameter_Sha
dowCopyType

TransferPrescriptParameter_Key=$TransferPrescriptParameter_Key
```

```
TransferPrescriptParameter_ForceBaseImage=$TransferPrescriptParameter_For
ceBaseImage

TransferPrescriptParameter_IsLogTruncation=$TransferPrescriptParameter_Is
LogTruncation

TransferPrescriptParameter_LatestEpochSeenByCore=$TransferPrescriptParame
ter_LatestEpochSeenByCore" > ~/PreTransferScriptResult.txt

exit 0
```

PostTransferScript.sh

The PostTransferScript is executed on the agent side after transferring a snapshot.

The following script stores the values from input parameters in the Pre(Post)TransferScriptResult.txt which is located and stored in the root home directory.

Sample PostTransferScript

```
#!/bin/bash

echo
"TransferPostscriptParameter_VolumeNames=$TransferPostscriptParameter_Vol
umeNames

TransferPostscriptParameter_ShadowCopyType=$TransferPostscriptParameter_S
hadowCopyType

TransferPostscriptParameter_Key=$TransferPostscriptParameter_Key

TransferPostscriptParameter_ForceBaseImage=$TransferPostscriptParameter_F
orceBaseImage

TransferPostscriptParameter_IsLogTruncation=$TransferPostscriptParameter_
IsLogTruncation

TransferPostscriptParameter_LatestEpochSeenByCore=$TransferPostscriptPara
meter_LatestEpochSeenByCore" > ~/PostTransferScriptResult.txt

exit 0
```

This page is intentionally left blank.

Glossary

Agent

An agent is a machine or server that is protected or to be protected by AppAssure 5.

AppAssure 5

AppAssure 5 sets a new standard for unified data protection by combining backup, replication, and recovery in a single solution that is engineered to be the fastest and most reliable backup for protecting virtual machines (VM), as well as physical and cloud environments.

Central Management Console

The AppAssure 5 Central Management Console is a multi-core management portal. It simplifies the process of managing multiple deployments of the AppAssure 5 Core. Using the Central Management Console, you can group and manage the deployments through a single, Web-based interface.

Checksum

A checksum is a function that creates blocks of data that are used for the purpose of detecting accidental errors that are created during transmission or storage.

Cluster

See "Windows Failover Cluster" on page 247.

Cluster Continuous Replication (CCR)

A non-shared storage failover cluster solution, that uses built-in asynchronous log shipping technology to create and maintain a copy of each storage group on a second server in a failover cluster. CCR is designed to be either a one or two data center solution, providing both high availability and site resilience. It is one of two types of clustered mailbox server (CMS) deployments available in Exchange 2007.

Cluster Node

An individual machine that is part of a Windows Failover cluster.

Compression

The Storage Networking Industry Association (SNIA) defines compression as the process of encoding data to reduce its size.

Core

The AppAssure 5 Core is the central component of the AppAssure architecture. The Core provides the essential services for backup, recovery, retention, replication, archival, and management. In the context of replication, the Core is also called a source core. The source core is the originating core, while the target core is the destination.

Database Availability Group (DAG)

A set of up to 16 Microsoft Exchange Server 2010 Mailbox servers that provide automatic, database-level recovery from a database, server, or network failure. DAGs use continuous replication and a subset of Windows failover clustering technologies to provide high availability and site resilience. Mailbox servers in a DAG monitor each other for failures. When a Mailbox server is added to a DAG, it works with the other servers in the DAG to provide automatic, database-level recovery from database failures.

Encryption

Data is encrypted with the intent that it is only accessible to authorized users who have the appropriate decryption key. Data is encrypted using 256-bit AES in Cipher Block Chaining (CBC) mode. In CBC, each block of data is XORed with the previous ciphertext block before being encrypted, this way each new ciphertext block depends on all preceding plaintext blocks. A passphrase is used as an initialization vector.

Global Deduplication

The Storage Networking Industry Association (SNIA) defines data deduplication as the replacement of multiple copies of data—at variable levels of granularity—with references to a shared copy to save storage space or bandwidth. The AppAssure Volume Manager performs global data deduplication within a logical volume. The granularity level of deduplication is 8 KB. The scope of deduplication in AppAssure is limited to protected machines using the same repository and encryption key.

License Key

The license key, which is obtained when you register on AppAssure 5 License Portal for an account is used access the license portal. From the AppAssure 5 License Portal, you can download the AppAssure 5 Core and Agents, manage licenses and groups, track group activity, register machines, create accounts, invite users, and generate reports.

License Portal

The AppAssure 5 License Portal is a Web interface where users and partners register, download, activate, and manage AppAssure 5 licenses.

Live Recovery

AppAssure Live Recovery is an instant recovery technology for VMs and servers. It provides near-continuous access to data volumes in a virtual or physical server, letting you recover an entire volume with near-zero RTO and a RPO of minutes.

Local Console

The Local Console is a Web-based interface that lets you fully manage the AppAssure 5 Core.

Local Mount Utility

The Local Mount Utility (LMU) is a downloadable application that lets you mount a recovery point on a remote AppAssure 5 Core from any machine.

Log Truncation

Log truncation is a function that is used to save space by removing log records from the transaction log.

Machine

A machine, sometimes referred to as an agent, is a physical or virtual machine or server that is protected by the AppAssure 5 Core. In the context of replication, a core may also be referred to as a source core.

Management Roles

The AppAssure 5 Central Management Console introduces a new concept of management roles which lets you divide administrative responsibility among trusted data and service administrators as well as access control to support secure and efficient delegation of administration.

Mountability

Exchange mountability is a corruption detection feature that alerts administrators of potential failures and ensures that all data on the Exchange servers is recovered successfully in the event of a failure.

Object File System

The AppAssure Scalable Object Store is an object file system component. It treats all data blocks, from which snapshots are derived, as objects. It stores, retrieves, maintains, and replicates these objects. It is designed to deliver scalable input and output (I/O) performance in tandem with global data deduplication, encryption, and retention management. The Object File System interfaces directly with industry standard storage technologies.

Passphrase

A passphrase is a key used in the encryption the data. If the passphrase is lost, data cannot be recovered.

PowerShell Scripting

Windows PowerShell is a Microsoft .NET Framework-connected environment designed for administrative automation. AppAssure 5 includes comprehensive client SDKs for PowerShell scripting that enables administrators to automate the administration and management of AppAssure 5 resources by the execution of commands either directly or through scripts.

Quorum

For a failover cluster, the number of elements that must be online for a given cluster to continue running. The elements relevant in this context are cluster nodes. This term can also refer to the quorum-capable resource selected to maintain the configuration data necessary to recover the cluster. This data contains details of all of the changes that have been applied to the cluster database. The quorum resource is generally accessible to other cluster resources so that any cluster node has access to the most recent database changes. By default there is only one quorum resource per server cluster. A particular quorum configuration (settings for a failover cluster) determines the point at which too many failures stop the cluster from running.

Recovery Assure

Recovery Assure technology is used to perform automated recovery testing and verification of backups. It supports various file systems and servers.

Recovery Points

Recovery points are a collection of snapshots of various disk volumes. For example, C:, D:, and E.

Remote Core

A remote core represents an AppAssure 5 Core that is accessed by a non-core machine by way of the Local Mount Utility.

Replication

Replication is self-optimizing with a unique read-match-write (RMW) algorithm that is tightly coupled with deduplication. It represents the relationship between the target and source cores in the same site or across two sites with slow link in which the source core asynchronously transmits the data to the target or source core on a per agent basis.

Repository

A repository, which is managed by the AppAssure 5 Core, is a folder used to store snapshots that are captured from the protected servers and machines. The repository can reside on different storage technologies such as Storage Area Network (SAN), Direct Attached Storage (DAS), or Network Attached Storage (NAS).

Retention

Retention defines the length of time the backup snapshots of protected machines are stored on the AppAssure 5 Core. Retention policy is enforced on the recovery points through the rollup process.

Rollback

Rollback is the process of restoring volumes on a machine from recovery points.

Rollup

The rollup process is an internal nightly maintenance procedure that enforces the retention policy by collapsing and eliminating dated recovery points. AppAssure 5 reduces rollup to metadata operations only.

Seeding

In replication, the initial transfer of deduplicated base images and incremental snapshots of protected agents, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target core using external media, which is useful for large sets of data or sites with slow links.

Server Cluster

See "Windows Failover Cluster" on page 247.

Single Copy Cluster

A shared storage failover cluster solution, that uses a single copy of a storage group on storage that is shared between the nodes in the cluster. It is one of two types of clustered mailbox server deployments available in Exchange 2007.

Smart Agent

The AppAssure 5 Smart Agent is installed on the machines protected by the AppAssure 5 Core. The smart agent tracks the changed blocks on the disk volume and snapshots the changed blocks at a predefined interval of protection

Snapshot

A snapshot is a common industry term that defines the ability to capture and store the state of a disk volume at a given point, while applications are running. The snapshot is critical if system recovery is needed due to an outage or system failure. AppAssure snapshots are application aware, which means that all open transactions and rolling transaction logs are completed and caches are flushed prior to creating the snapshot. AppAssure uses Microsoft Volume Shadow Services (VSS) to facilitate application crash consistent snapshots.

SQL Attachability

SQL attachability is a test run within the AppAssure 5 Core to ensure that all SQL recovery points are without error and are available for backup in the event of a failure.

Target Core

The target core, which is sometimes referred to as replica core, is the AppAssure 5 Core receiving the replicated data from the source core.

Target Replica Machine

The instance of a protected machine on a target core is known as the target agent or replica agent.

Transport Layer Security

Transport Layer Security (TLS) is a modern cryptographic network protocol designed to ensure communication security over the Internet. This protocol, defined by the Internet Engineering Task Force, is the successor to Secure Sockets Layer (SSL). The SSL term is still generally used, and the protocols are interoperable (a TLS client can downgrade to communicate to an SSL server).

True Scale™

True Scale is the scalable architecture of AppAssure 5.

Universal Recovery

AppAssure Universal Recovery technology provides unlimited machine restoration flexibility. It enables you to perform monolithic recovery to- and from- any physical or virtual platform of your choice as well as incremental recovery updates to virtual machines from any physical or virtual source. It also lets you perform application-level, item-level, and object-level recovery of individual files, folders, email, calendar items, databases, and applications.

Virtual Standby

Virtual Standby is a physical-to-virtual (P2V) process that creates a clone virtual machine of a protected machine or agent. A Virtual Standby can be created using an ad-hoc or a continuous update export process. A Virtual Standby created using a continuous update is incrementally updated after every snapshot captured from the source agent.

Volume Manager

The AppAssure 5 Volume Manager manages objects and then stores and presents them as a logical volume. It leverages dynamic pipeline architecture to deliver TruScale™ scalability, parallelism, and asynchronous input-and-output (I/O) model for high throughput with minimal I/O latency.

White Labeling

AppAssure 5 provides the ability for providers of backup and disaster recovery services to white label or re-brand AppAssure 5 with their own identity; and then sell or distribute it as their own product or service.

Windows Failover Cluster

A group of independent computers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to provide service (a process known as failover). Users experience a minimum of disruptions in service. AppAssure 5 supports the protection of a number of SQL Server and Exchange Server cluster types.

This page intentionally left blank.

Index

A

B

C

D

E

F

R

S

T

U

V

W