

Dell DL4000 Appliance User's Guide



Notes, cautions, and warnings



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



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



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

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

2015 - 12

Rev. A01

Contents

1 Introduction to Dell DL4000 Appliance.....	10
Core technologies.....	10
Live Recovery.....	11
Verified Recovery.....	11
Universal Recovery.....	11
True Global Deduplication.....	11
True Scale architecture.....	11
Deployment architecture.....	12
Smart Agent.....	14
DL4000 Core.....	14
Snapshot process.....	14
Replication of disaster recovery site or service provider.....	15
Recovery.....	15
Product features	15
Repository.....	16
True Global Deduplication	16
Encryption.....	17
Replication.....	18
Recovery-as-a-Service (RaaS).....	19
Retention and archiving.....	19
Virtualization and cloud.....	20
Alerts and event management.....	20
License portal.....	20
Web console.....	20
Service management APIs.....	21
2 Working with the DL4000 Core.....	22
Accessing the DL4000 Core Console.....	22
Updating trusted sites in Internet Explorer.....	22
Configuring browsers to remotely access the Core Console.....	22
Roadmap for configuring the Core	23
Managing licenses	24
Changing a license key	24
Contacting the license portal server	24
Changing the AppAssure language manually.....	25
Changing the OS language during installation.....	25
Managing Core settings	26
Changing the Core display name	26

Adjusting the nightly job time	26
Modifying the transfer queue settings	26
Adjusting the client time-out settings	27
Configuring deduplication cache settings	27
Modifying engine settings	28
Modifying database connection settings	29
About repositories	29
Roadmap for managing a repository	30
Creating a repository	30
Viewing repository details.....	33
Modifying repository settings	33
Expanding an existing repository.....	34
Adding a storage location to an existing repository	34
Checking a repository	36
Deleting a repository	36
Remounting volumes.....	36
Recovering a repository.....	37
Managing security	37
Adding an encryption key	38
Editing an encryption key	38
Changing an encryption key passphrase	38
Importing an encryption key	39
Exporting an encryption key	39
Removing an encryption key	39
Managing cloud accounts	39
Adding a cloud account.....	40
Editing a cloud account.....	41
Configuring cloud account settings.....	41
Understanding replication	42
About protecting workstations and servers	42
About replication	42
About seeding	44
About failover and failback	45
About replication and encrypted recovery points	45
About retention policies for replication	45
Performance considerations for replicated data transfer	45
Roadmap for performing replication	46
Replicating to a self-managed core.....	47
Replicating to a core managed by a third party.....	50
Monitoring replication	53
Managing replication settings	54
Removing replication	54

Removing a protected machine from replication on the source Core.....	55
Removing a protected machine on the target Core.....	55
Removing a target Core from replication.....	55
Removing a source Core from replication.....	55
Recovering replicated data	56
Roadmap for failover and failback	56
Setting up an environment for failover	56
Performing failover on the target Core	57
Performing failback	57
Managing events	58
Configuring notification groups	59
Configuring an email server and email notification template	60
Configuring repetition reduction	61
Configuring event retention	61
Managing recovery	62
About system information	62
Viewing system information	62
Downloading installers	62
About the agent installer	62
Downloading and installing the agent installer	63
About the local mount utility	63
Downloading and installing the local mount utility	63
Adding a core to the local mount utility	64
Mounting a recovery point by using the local mount utility	65
Dismounting a recovery point by using the local mount utility	66
About the local mount utility tray menu	66
Using Core and agent options.....	66
Managing retention policies	67
Archiving to a cloud.....	67
About archiving	67
Creating an archive	68
Setting a scheduled archive	68
Pausing or resuming scheduled archive	69
Editing a scheduled archive	70
Checking an archive	71
Importing an archive	71
Managing SQL attachability	72
Configuring SQL attachability settings	72
Configuring nightly SQL attachability checks and log truncation	73
Managing exchange database mountability checks and log truncation	73
Configuring exchange database mountability and log truncation	74
Forcing a mountability check	74

Forcing checksum checks	74
Forcing log truncation	75
Recovery point status indicators	75
3 Managing Your Appliance.....	77
Monitoring the status of the Appliance.....	77
Provisioning storage.....	77
Provisioning selected storage.....	78
Deleting space allocation for a virtual disk.....	79
Resolving failed tasks.....	79
Upgrading your Appliance.....	79
Repairing your Appliance.....	80
4 Protecting workstations and servers.....	81
About protecting workstations and servers	81
Configuring machine settings	81
Viewing and modifying configuration settings	81
Viewing system information for a machine	82
Configuring notification groups for system events	82
Editing notification groups for system events	84
Customizing retention policy settings	86
Viewing license information	88
Modifying protection schedules	88
Modifying transfer settings	89
Restarting a service	91
Viewing machine logs	92
Protecting a machine	92
Deploying the agent software when protecting an agent.....	94
Creating custom schedules for volumes	95
Modifying exchange server settings	95
Modifying SQL server settings	96
Deploying an agent (push install)	96
Replicating a new agent	97
Managing machines	98
Removing a machine	98
Replicating agent data on a machine	98
Setting replication priority for an agent	99
Canceling operations on a machine	99
Viewing machine status and other details	100
Managing multiple machines	101
Deploying to multiple machines	101
Monitoring the deployment of multiple machines	105

Protecting multiple machines	105
Monitoring the protection of multiple machines	107
Managing snapshots and recovery points	107
Viewing recovery points	108
Viewing a specific recovery point.....	108
Mounting a recovery point for a Windows machine	109
Dismounting select recovery points.....	110
Dismounting all recovery points.....	110
Mounting a recovery point volume on a Linux machine	110
Removing recovery points	111
Deleting an orphaned recovery point chain.....	111
Forcing a snapshot	112
Pausing and resuming protection	112
Restoring data	113
Backup.....	113
About exporting protected data from Windows machines to virtual machines.....	114
Exporting backup information from your Microsoft Windows machine to a virtual machine	116
Exporting Windows data using ESXi export	116
Exporting Windows data using VMware workstation export	118
Exporting Windows data using Hyper-V export	120
Exporting Microsoft Windows data using Oracle VirtualBox export	123
Virtual Machine Management.....	125
Performing a rollback	129
Performing a rollback for a Linux machine by using the command line.....	130
About bare metal restore for Windows machines	131
Prerequisites for performing a bare metal restore for a Windows machine	131
Roadmap for performing a bare metal restore for a Windows machine	132
Creating a bootable CD ISO image.....	132
Loading a boot CD.....	134
Launching a restore from the Core	135
Mapping volumes	135
Viewing the recovery progress	136
Starting the restored target server	136
Repairing startup problems.....	136
Performing a bare metal restore for a Linux machine	137
Installing the screen utility.....	138
Creating bootable partitions on a Linux machine.....	138
Viewing events and alerts	139
5 Protecting server clusters.....	140
About server cluster protection	140

Supported applications and cluster types	140
Protecting a cluster	141
Protecting nodes in a cluster	142
Process of modifying cluster node settings	143
Roadmap for configuring cluster settings	143
Modifying cluster settings	144
Configuring cluster event notifications	144
Modifying the cluster retention policy	145
Modifying cluster protection schedules	146
Modifying cluster transfer settings	146
Converting a protected cluster node to an agent	147
Viewing server cluster information	147
Viewing cluster system information	147
Viewing summary information	148
Working with cluster recovery points	148
Managing snapshots for a cluster	148
Forcing a snapshot for a cluster	149
Pausing and resuming cluster snapshots	149
Dismounting local recovery points	149
Performing a rollback for clusters and cluster nodes	150
Performing a rollback for CCR (Exchange) and DAG clusters	150
Performing a rollback for SCC (Exchange, SQL) clusters.....	150
Replicating cluster data	150
Removing a cluster from protection	150
Removing cluster nodes from protection	151
Removing all nodes in a cluster from protection	151
Viewing a cluster or node report	152
6 Reporting.....	153
About reports	153
About the reports toolbar	153
About compliance reports	153
About errors reports	154
About the Core Summary Report	154
Repositories summary	154
Agents summary	155
Generating a report for a Core or agent	155
About the Central Management Console Core reports	156
Generating a report from the Central Management Console	156
7 Completing a full recovery of the DL4000 Appliance.....	157
Creating a RAID 1 partition for the operating system.....	157

Installing the operating system.....	158
Running the recovery and update utility.....	158
8 Changing the host name manually.....	160
Stopping the Core service.....	160
Deleting server certificates.....	160
Deleting Core server and registry keys.....	160
Launching the Core with the new host name.....	161
Changing the display name	161
Updating trusted sites in Internet Explorer.....	161
9 Appendix A— scripting.....	162
About powershell scripting	162
Powershell scripting prerequisites	162
Testing scripts	162
Input parameters	163
VolumeNameCollection (namespace Replay.Common.Contracts.Metadata.Storage)	167
Pretransferscript.ps1	168
Posttransferscript.ps1	168
Preexportscript.ps1	169
Postexportscript.ps1	170
Prenightlyjobscript.ps1	170
Postnightlyjobscript.ps1.....	172
Sample scripts	174
10 Getting help.....	175
Finding documentation and software updates.....	175
Contacting Dell.....	175
Documentation feedback.....	175

Introduction to Dell DL4000 Appliance

This chapter provides an introduction and overview of DL4000. It describes the features, functionality, and architecture, and consists of the following topics:

- [Core technologies](#)
- [True Scale architecture](#)
- [Deployment architecture](#)
- [Product features](#)

Your appliance 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.

Your appliance is capable of handling up to petabytes of data 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 (DR) and compliance.

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

Your appliance combines the following technologies:

- [Live Recovery](#)
- [Verified Recovery](#)
- [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, your appliance 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.

AppAssure addresses the complexity and inefficiency of legacy tools through its 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 offers these technological advances while dramatically reducing IT management and storage costs.

Core technologies

Details about the core technologies of AppAssure 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.

The 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 Recovery Time Objective (RTO) and Recovery Point Objective (RPO) service-level agreements.

Verified Recovery

Verified Recovery enables you to perform automated recovery testing and verification of backups. It includes, but is not limited to, file systems: - Microsoft Exchange 2007, 2010, and 2013, and different versions of Microsoft SQL Server 2005, 2008, 2008 R2, 2012 and 2014. Verified Recovery provides recoverability of applications and backups in virtual and physical environments. It 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 systems to virtual machines, virtual machines to virtual machines, virtual machines to physical systems, or physical systems to physical systems, and carry out bare metal restores to dissimilar hardware. For example, P2V, V2V, V2P, P2P, P2C, V2C, C2P, and C2V.

Universal Recovery technology also accelerates cross-platform moves among virtual machines. For example, moving from VMware to Hyper-V or Hyper-V to VMware. It builds in application-level, item-level, and object-level recovery (individual files, folders, e-mail, calendar items, databases, and applications). With AppAssure, you can recover or export physical to cloud, or virtual to cloud.

True Global Deduplication

Your appliance provides true global deduplication that reduces your physical disk drive capacity requirements by offering space reduction ratios exceeding 50:1, while still meeting the data storage requirements. AppAssure 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.

True Scale architecture

Your appliance is built on AppAssure True Scale architecture. It leverages dynamic, multi-core pipeline architecture that is optimized to consistently deliver solid performance for your enterprise environments. True Scale is designed from the ground up to linearly scale and efficiently store and manage big data, and deliver RTOs and RPOs of minutes without compromising performance. It comprises of a purpose-built

object and a volume manager with integrated global deduplication, compression, encryption, replication, and retention. The following diagram describes the AppAssure True Scale architecture.




Figure 1. AppAssure True Scale architecture

The AppAssure Volume Manager and Scalable Object Store serve as the foundation of the AppAssure True Scale architecture. The scalable object store stores block-level snapshots that are captured from virtual and physical servers. The volume manager manages the numerous object stores by providing a common repository or just-in-time storage for only what is needed. The Object Store concurrently supports everything with asynchronous I/O that delivers high throughput with minimal latency and maximizes system utilization. The repository resides on different storage technologies such as Storage Area Network (SAN), Direct Attached Storage (DAS), or Network Attached Storage (NAS).

The role of the AppAssure Volume Manager is similar to the role of the volume manager in an operating system. It takes various storage devices which can be of different sizes and types and combines them into logical volumes, using striped or sequential allocation policies. The object store saves, retrieves, maintains, and then replicates objects that are derived from application-aware snapshots. The volume manager delivers scalable I/O performance in tandem with global data deduplication, encryption, and retention management.

Deployment architecture

Your appliance is a scalable backup and recovery product that is flexibly deployed within the enterprise or as a service delivered by a managed service provider. The type of deployment depends on the size and requirements of the customer. Preparing to deploy your appliance involves planning the network storage topology, core hardware and disaster recovery infrastructure, and security.

The deployment architecture consists of local and remote components. The remote components may be optional for those environments that do not require leveraging a disaster recovery site or a managed service provider for off-site recovery. A basic local deployment consists of a backup server called the Core and one or more protected machines. The off-site component is enabled using replication that provides full recovery capabilities in the DR site. The Core uses base images and incremental snapshots to compile recovery points of protected machines.

Additionally, your appliance is application-aware because it can detect the presence of Microsoft Exchange and SQL and their respective databases and log files, and then automatically group these

volumes with dependency for comprehensive protection and effective recovery. This ensures that you never have incomplete backups when you are performing recoveries. Backups are performed by using application-aware block-level snapshots. Your appliance can also perform log truncation of the protected Microsoft Exchange and SQL servers.

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




Figure 2. Basic deployment architecture

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




Figure 3. Multi-Core deployment architecture

Smart Agent

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

The Smart Agent is application-aware and is dormant when not in use, with near zero (0) percent CPU utilization and less than 20 MB of memory overhead. When the Smart Agent is active, it uses up to 2 to 4 percent processor utilization and less than 150 MB memory, which includes transferring the snapshots to the Core.

The Smart Agent is application-aware and it detects the type of application that is installed and also the location of the data. It automatically groups data volumes with dependency, such as databases, and then logs them together for effective protection and rapid recovery. After the AppAssure Agent software is configured, it uses smart technology to keep track of changed blocks on the protected disk volumes. When the snapshot is ready, it is rapidly transferred to the Core using intelligent multi-threaded, socket-based connections. To preserve CPU bandwidth and memory on the protected machines, the smart agent does not encrypt or deduplicate the data at the source and protected machines are paired with a Core for protection.

DL4000 Core

The Core is the central component of the deployment architecture. The Core stores and manages all of the machine backups and provides core services for backup, recovery, and retention; replication, archival, and management. The Core is a self-contained network-addressable computer that runs a 64-bit of Microsoft Windows operating system. Your appliance performs target-based inline compression, encryption, and deduplication of the data received from the protected machine. The Core then stores the snapshot backups in repositories such as, Storage Area Network (SAN) or Direct Attached Storage (DAS).

The repository can also reside on internal storage within the Core. The Core is managed by accessing the following URL from a Web browser: **<https://CORENAME:8006/apprecovery/admin>**. Internally, all core services are accessible through REST APIs. The Core services can be accessed from within the core or directly over the Internet from any application that can send an HTTP/HTTPS request and receive an HTTP/HTTPS response. All API operations are performed over SSL and mutually authenticated using X.509 v3 certificates.

Cores are paired with other cores for replication.

Snapshot process

A snapshot is when a base image is transferred from a protected machine to the Core. This is the only time a full copy of the machine is transported across the network under normal operation, followed by incremental snapshots. AppAssure Agent software for Windows uses Microsoft Volume Shadow copy Service (VSS) to freeze and quiesce application data to disk to capture a file-system-consistent and an application-consistent backup. When a snapshot is created, the VSS, and the writer on the target server prevent content from being written to the disk. When the writing of content to disk is halted, all disk I/O operations are queued and resume only after the snapshot is complete, while the operations already in flight are completed and all open files are closed. The process of creating a shadow copy does not significantly impact the performance of the production system.

AppAssure uses Microsoft VSS because it has built-in support for all Windows internal technologies such as NTFS, Registry, Active Directory, to flush data to disk before the snapshot. Additionally, other enterprise applications, such as Microsoft Exchange and SQL, use VSS Writer plug-ins to get notified when a snapshot is being prepared and when they have to flush their used database pages to disk to bring the database to a consistent transactional state. It is important to note that VSS is used to quiesce system and application data to disk; it is not used to create the snapshot. The captured data is immediately transferred and stored on the Core. Using VSS for backup does not render the application server in backup mode for an extended period of time because the time taken to create the snapshot is seconds and not hours. Another benefit of using VSS for backups is that it lets the AppAssure Agent software to take a snapshot of large quantities of data at one time because the snapshot works at the volume level.

Replication of disaster recovery site or service provider

The replication process requires a paired source-target relationship between two cores. The source core copies the recovery points of the protected machines 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 workflows that let you request connections and receive automatic feedback notifications. For the initial transfer of data, you can perform data seeding using external media, which is useful for large sets of data or sites with slow links.

In the case of a severe outage, your appliance supports failover and failback in replicated environments. In case of a comprehensive outage, the target core in the secondary site can recover instances from replicated protected machines and immediately commence protection on the failed-over machines. After the primary site is restored, the replicated core can fail-back data from the recovered instances back to protected machines at the primary site.

Recovery

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

Product features

You can manage protection and recovery of critical data using the following features and functionality:

- [Repository](#)
- [True Global Deduplication \(Features\)](#)
- [Encryption](#)
- [Replication](#)
- [Recovery-as-a-Service \(RaaS\)](#)
- [Retention and archiving](#)
- [Virtualization And Cloud](#)
- [Alerts and Event Management](#)
- [License portal](#)
- [Web console](#)
- [Service Management APIs](#)

Repository

The 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 allows you to 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 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 your appliance performs incremental block-level backups of the machine. The 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 4. Diagram of deduplication

Your appliance performs target-based inline data deduplication, where 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 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 your appliance 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 50:1.

Encryption

Your appliance provides integrated encryption to protect backups and data-at-rest from unauthorized access and use, ensuring data privacy. Only a 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. 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 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's 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, your appliance 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 from an AppAssure core and transmitting them to another AppAssure core in a separate location for the purpose of disaster recovery. The process requires a paired source-target relationship between two or more cores.

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

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




Figure 5. Basic replication architecture

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. Seeding is the initial transfer of deduplicated base images and incremental snapshots of the protected machines. The data can add up to hundreds or thousands of gigabytes. Initial replication can be seeded to the target core using external media. This is 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 external media, the archive can span across multiple devices. During the seeding process, the incremental recovery points

replicate to the target site. After data has been transferred to the target core, the newly replicated incremental recovery points automatically synchronize.

Recovery-as-a-Service (RaaS)

Managed service providers (MSPs) can fully leverage the appliance 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 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

In your appliance, backup and retention policies are flexible and, therefore, easily configurable. The ability to tailor retention policies to the needs of an organization not only helps to meet compliance requirements, but does so without compromising on 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, it can also be used for seeding replication data to a target core.




Figure 6. Custom retention policy

In your appliance, retention policies can be customized to specify the length of time a backup recovery point is maintained. As the age of the recovery points approaches the end of their retention period, the recovery points 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 grows rapidly. Your

appliance 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. As 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, days, weeks, months, and years. The retention policy is enforced by the nightly rollup process.

For long-term archiving, your appliance provides the ability to 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 spans across multiple devices based on the available space on the media. The archive also can be locked with a passphrase. 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 Core is cloud-ready, which allows you to leverage the compute capacity of the cloud for recovery.

Your appliance can export any protected or replicated machine to a virtual machine, such as licensed versions of VMware or Hyper-V. You can perform a one-time virtual export, or you can establish a virtual standby VM by establishing a continuous virtual export. 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 virtual machine export types are VMware Workstation/Server on a folder; direct export to a vSphere/VMware ESX(i) host; export to Oracle VirtualBox; and export to Microsoft Hyper-V Server on Windows Server 2008 (x64), 2008 R2, 2012 (x64), and 2012 R2 (including support for Hyper-V generation 2 VMs)

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

Alerts and event management

In addition to HTTP REST API, your appliance also includes an extensive set of features for event logging and notification using e-mail, 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 environment. In Windows-only environments, only the Windows Event Log is used.

License portal

The 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

Your appliance features a new web-based central console that manages distributed cores from one central location. MSPs and enterprise customers with multiple distributed cores can deploy the central console to get a unified view for central management. The central console provides the ability to

organize the managed cores in hierarchical organizational units. These organizational units can represent business units, locations, or customers for MSPs with role-based access. The central console can also run reports across managed cores.

Service management APIs

Your appliance comes bundled with a service management API and provides programmatic access to all of the functionality available through the 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. This approach facilitates easy integration with any web application such as relationship management methodology (RMM) tools or billing systems. Also included is an SDK client for PowerShell scripting.

Working with the DL4000 Core

Accessing the DL4000 Core Console

To access the Core Console:

1. Update trusted sites in your browser. See [Updating Trusted Sites In Internet Explorer](#).
2. Configure your browsers to remotely access the Core Console. See [Configuring Browsers To Remotely Access The Core Console](#).
3. Perform one of the following to access the Core Console:
 - Log on locally to your DL4000 core server, and then double-click the **Core Console** icon.
 - Type one of the following URLs in your web browser:
 - **https://<yourCoreServerName>:8006/apprecovery/admin/core**
 - **https://<yourCoreServerIPAddress>:8006/apprecovery/admin/core**

Updating trusted sites in Internet Explorer

To update trusted sites in Microsoft Internet Explorer:

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

Configuring browsers to remotely access the Core Console

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



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



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



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

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

Configuring browser settings in Internet Explorer and Chrome

To modify browser settings in Internet Explorer and Chrome:

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

Configuring Mozilla Firefox browser settings



NOTE: To modify Mozilla Firefox browser settings in the latest versions of Firefox, disable protection. Right-click the Site Identify button (located to the left of the URL), go to **Options** and click on **Disable protection for now**.

To modify Mozilla Firefox browser settings:

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

Roadmap for configuring the 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 Core, you can then protect agents and perform recovery.

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

- Create a repository
- Configure encryption keys
- Configure event notification
- Configure retention policy
- Configure SQL attachability

 **NOTE:** If you are using this Appliance, it is recommended that you use the **Appliance** tab to configure the Core. For more information about configuring the Core after initial installation, see the *Dell DL4000 Appliance Deployment Guide* at dell.com/support/home.

Managing licenses

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

The Licensing page includes the following information:

- License type
- License status
- License constraints
- 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

Changing a license key

To change a license key:

1. Navigate to the Core Console.
2. Select **Configuration** → **Licensing**.
The **Licensing** page appears.
3. From the **License Details** section, click **Change License**.
The **Change License** dialog box appears.
4. In the **Change License** dialog box, enter the new license key and then click **Continue**.

Contacting the license portal server

The 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.

To contact the portal server:

1. Navigate to the Core Console.
2. Click **Configuration** → **Licensing**.
3. From the **License Server** option, click **Contact Now**.

Changing the AppAssure language manually

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

To change the AppAssure language to the desired language:

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

Changing the OS language during installation

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

To change OS language:

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

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

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

9. Log out and log in again to Windows for the change to take effect.

Managing Core settings

The Core settings are used to define various 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

 **NOTE:** It is recommended that you select a permanent display name during the initial configuration of your Appliance. If you change it later, you must perform several steps manually to ensure that the new host name takes effect and the appliance functions properly. For more information, see [Changing The Host Name Manually](#).

To change the Core display name:

1. Navigate to the Core Console.
2. Click **Configuration** → **Settings**.
3. In the **General** pane, click **Change**.
The **General Settings** dialog box appears.
4. In the **Display Name** text box, enter a new display name for the Core.
This is the name that will display in the Core Console. You can enter up to 64 characters.
5. In the **Web Server Port** text box, enter a port number for the web server. The default is 8006.
6. In the **Service Port**, enter a port number for the service. The default is 8006.
7. Click **OK**.

Adjusting the nightly job time

To adjust the nightly job time:

1. Navigate to the Core Console.
2. Click **Configuration** → **Settings**.
3. In the **Nightly Jobs** area, click **Change**.
The **Nightly Jobs** dialog box appears.
4. In the **Nightly Jobs Time** text box, enter a new time to perform the nightly jobs.
5. 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.

To modify the transfer queue settings:

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

Adjusting the client time-out settings

To adjust the client time-out settings:

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

Configuring deduplication cache settings

To configure deduplication cache settings:

1. Navigate to the Core Console.
2. Click **Configuration** → **Settings**.
3. In the **Deduplication Cache Configuration** area, click **Change**.
The **Deduplication Cache Configuration** dialog box appears.
4. In the **Primary Cache Location** text box, enter an updated value to change the primary cache location.
5. In the **Secondary Cache Location** text box, enter an updated value to change the secondary cache location.
6. In the **Metadata Cache Location** text box, enter an updated value to change the metadata cache location.
7. In the **Dedupe Cache Size** text box, enter a value corresponding to the amount of space you want to allocate for the deduplication cache.
From the unit size drop-down field, select either GB (gigabytes) or TB (terabytes), to specify the unit of measurement for the value in the Dedupe Cache Size text box.

8. Click **OK**.



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

Modifying engine settings

To modify the engine settings:

1. Navigate to the Core Console.
2. Click **Configuration** → **Settings**
3. In the **Replay Engine Configuration** pane, click **Change**.
The **Replay Engine Configuration** dialog box appears.
4. Enter the configuration information described as follows:

Text Box	Description
IP address	• To use the preferred IP address from your TCP/IP, click Automatically Determined • To manually enter an IP address, click Use a specific address.
Preferable Port	Enter a port number or accept the default setting (8007 is the default port). The port is used to specify the communication channel for the engine.
Port in use	Represents the port that is in use for the Replay Engine configuration.
Allow port auto-assigning	Click for allow for automatic TCP port assignment.
Admin Group	Enter a new name for the administration group. The default name is BUILTIN Administrators .
Minimum 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.

5. Click **OK**.

Modifying database connection settings

To modify database connection settings:

1. Navigate to the Core Console.
2. Click **Configuration** → **Settings**
3. In the **Database Connection Settings** area, choose one of the following:
 - Click **Apply Default**.
 - Click **Change**.

The **Database Connection Settings** dialog box appears.

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

Text Box	Description
Host Name	Enter a host name for the database connection.
Port	Enter a port number for the database connection.
User Name (optional)	Enter a user name for accessing and managing the database connection settings. It is used to specify the log in credentials for accessing the database connection.
Password (optional)	Enter a password for accessing and managing the database connection settings.
Retain event and job history for, days	Enter the number of days to retain the event and job history for the database connection.
Max connection pool size	Sets the maximum number of database connections cached to allow dynamic reuse. Default setting is 100.
Min connection pool size	Sets the minimum number of database connections cached to allow dynamic reuse. Default setting is 0.

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

About repositories

A repository stores 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 Core preallocates 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:** DL4000 repositories are stored on primary storage devices. Archival storage devices such as Data Domain are not supported due to performance limitations. Similarly, repositories must 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

The roadmap for managing a repository covers tasks such as creating, configuring, and viewing a repository, and includes the following topics:

- [Accessing The Core Console](#)
- [Creating A Repository](#)
- [Viewing Repository Details](#)
- [Modifying Repository Settings](#)
- [Adding A Storage Location To An Existing Repository](#)
- [Checking A Repository](#)
- [Deleting A Repository](#)
- [Recovering A Repository](#)

 **NOTE:** It is recommended that you use the **Appliance** tab to configure repositories.

Before you begin using your appliance, you must set up one or more repositories on the 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 various tasks such as specifying where to locate the data storage on the Core server, how many locations can be added to each repository, the name of the repository, how many current operations the repositories support.

When you create a repository, the Core preallocates 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.

You can add or modify repositories in the Core Console.

Creating a repository

 **NOTE:** If you are using this appliance as a SAN, it is recommended that you use the **Appliance** tab to create repositories, see [Provisioning selected storage](#).

Perform the following to manually create a repository:

1. Navigate to the Core Console.
2. Click **Configuration** → **Repositories**.
3. Click **Add new**.
The **Add New Repository** dialog box appears.
4. 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 sequentially adds a number to the new repository starting with 1. You can change the name as needed. You can enter up to 150 characters.
Concurrent Operations	Define the number of concurrent requests that you want the repository to support. By default the value is 64.
Comments	Optionally, enter a descriptive note about this repository.

5. To define the specific storage location or volume for the repository, click **Add Storage Location**.



CAUTION: 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: Repositories are stored on primary storage devices. Archival storage devices such as Data Domain are not supported due to performance limitations. Similarly, repositories must 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 is displayed.

6. 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.
 - To specify a local machine, click **Add file on local disk**, and then enter the information described as follows:

Text Box	Description
Data Path	Enter the location for storing the protected data; for example, type X:\Repository\Data . 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.
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.

- Or, to specify a network share location, click **Add file on CIFS share**, and then enter the information described as follows:

Text Box	Description
UNC Path	Enter the path for the network share location.

Text Box	Description 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.
7. In the Details pane, click Show/Hide Details and enter the details for the storage location described as follows:	
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.  NOTE: 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 an NTFS volume using Windows 8 or Windows Server 2012, the file size limit is 256 TB.  NOTE: 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 the value is 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 controls the caching. If set to Sync, Windows controls the caching as well as the synchronous input/output.

Text Box	Description
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.

8. Click **Save**.
The **Repositories** screen is displayed to include the newly added storage location.
9. Repeat step 4 through step 7 to add more storage locations for the repository.
10. Click **Create** to create the repository.
The **Repository** information appears in the **Configuration** tab.

Viewing repository details

To view repository details:

1. Navigate to the Core Console.
2. Click **Configuration** → **Repositories**.
3. Click > next to the **Status** column of the repository for which you want to view the details.
4. From the expanded view, you can perform the following actions:
 - Modify Settings
 - Add a Storage Location
 - Check a Repository
 - Delete a Repository

Details also display for the repository and include the storage locations and statistics. Storage location details include the metadata path, data path, and size. Statistical information includes:

- Deduplication — Reported as the number of block dedupe hits, block dedupe misses, and block compression rate.
- Record I/O — Consisting of the rate (MB/s), read rate (MB/s), and write write (MB/s).
- Storage Engine — Include the rate (MB/s) read rate (MB/s), and write write (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 create a new storage location for the repository.

To modify repository settings:

1. Navigate to the Core Console.
2. Click **Configuration** → **Repositories**.
3. Click the Settings icon next to the Compression Ratio column below the **Actions** button, and then **Settings**.
The **Repository Settings** dialog box appears.
4. Edit the repository information described as follows:

Field	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 that you want the repository to support.
Enable Deduplication	To turn off deduplication, clear this check box. To enable deduplication, select this check box.  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, retains the deduplication values in place at the time the data was captured from the protected machine.
Enable Compression	To turn off compression, clear this check box. To enable compression, select this check box.  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, retains the compression values in place at the time the data was captured from the protected machine.

5. Click **Save**.

Expanding an existing repository

If you add another MD1200 DAS to your appliance, you can use the available storage to expand an existing repository.

To expand an existing repository:

1. After you install the MD1200 DAS, open the Core Console and select the **Appliance** tab, click **Tasks**.
2. On the **Tasks** screen, next to the new storage, click **Provision**.
3. On the **Provisioning Storage** screen, select **Expand the existing repository**, select the repository that you want to expand.
4. Click **Provision**.
The **Tasks** screen displays the **Status Description** next to the storage device as **Provisioned**.

Adding a storage location to an existing repository

Adding a storage location lets you define where you want to store the repository or volume.

To add a storage location to an existing repository:

1. Click > next to the **Status** column of the repository for which you want to add a storage location.
2. Click **Add Storage Location**.
The **Add Storage Location** dialog box appears.

3. Specify how to add the file for the storage location. You can choose to add the file on the local disk or on a CIFS share.

- To specify a local machine, click **Add file on local disk**, enter the information as follows:

Text Box	Description
Metadata Path	Enter the location for storing the protected metadata.
Data Path	Enter the location for storing the protected data.

- To specify a network share location, click **Add file on CIFS share**, enter the information as follows:

Text Box	Description
UNC Path	Enter the path for the network share location.
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** section, click **Show/Hide Details** and enter the details for the storage location described as follows:

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.  NOTE: If the storage location is an NTFS volume using Windows XP or Window 7, the file size limit is 16 TB. If the storage location is an NTFS volume using Windows 8 or Windows Server 2012, the file size limit is 256 TB.  NOTE: To validate the operating system, 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; however the recommended setting is Off. If set to Off, AppAssure controls the caching.

Text Box	Description
	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.

- Click **Save**.
The **Repositories** screen is displayed to include the newly added storage location.
- Repeat Step 4 through Step 7 to add more storage locations for the repository.
- Click **OK**.

Checking a repository

The appliance can perform a diagnostic check of a repository volume when errors occur. Core errors could be the result of it being improperly shut down, or a hardware failure, among other reasons.

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

To check a repository:

- On the **Configuration** tab, click **Repositories**, select > next to the repository that you want to check.
- In the **Actions** pane, click **Check**.
The **Check Repository** dialog box appears.
- In the **Check Repository** dialog box, click **Check**.

 **NOTE:** If the check fails, restore the repository from an archive.

Deleting a repository

To delete a repository:

- On the **Configuration** tab, click **Repositories**, select > next to the repository that you want to delete.
- In the **Actions** pane, click **Delete**.
- In the **Delete Repository** dialog box, click **Delete**.

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

When you delete a repository, then you must go through the Open Manage System Administrator and delete the virtual disks that housed the repository. After you delete the virtual disks, you can re-provision the disks and recreate the repository.

Remounting volumes

To remount the volumes:

- Navigate to the Core Console.
- Appliance** → **Tasks**.
- Click **Remount Volumes**.

The Volumes remount.

Resolving foreign volumes

If a provisioned MD1200 is powered off or disconnected and then later powered back on, an event appears on the Core Console reporting that the MD1200 is connected. However, no task appears on the **Appliance** tab **Tasks** screen that permits you to recover it. The **Enclosures** screen reports the MD1200 as being in a foreign state and the repositories on the foreign virtual disks as off-line.

To resolve foreign volumes:

1. From the Core Console, select the **Appliance** tab, and then click **Remount Volumes**.
The volumes remount.
2. Select the **Configuration** tab, and then click **Repositories**.
3. Expand the repository with the red status indicator by clicking > next to **Status**.
4. To verify the repository integrity, under **Actions**, click **Check**.

Recovering a repository

When the appliance fails to import a repository, it reports the failure in the **Tasks** screen with the task status indicated by a red circle, and the status description reporting **Error, Completed — Exception**. To view the error details from the **Tasks** screen, expand the task details by clicking > next to the **Status** column. **Status Details** reports that the recovery task status is exception, and the **Error Message** column provides additional details about the error condition.

To recover a repository from a failed import state:

1. Navigate to the Core Console.
The **Repositories** screen displays the failed repository with a red status indicator.
2. Click **Configuration** → **Repositories**.
3. Expand the failed repository by clicking > next to **Status**.
4. From the **Actions** section, click **Check**, and then click **Yes** to confirm that you want to run the check.
The appliance recovers the repository.

Managing security

The Core can encrypt protected machine snapshot data within the repository. Instead of encrypting the entire repository, you can specify an encryption key during the protection of a machine in a repository which lets the keys be reused for different protected machines. Encryption does not affect performance, as each active encryption key creates an encryption domain, thus letting a single core support multitenancy 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 Core.
- There is no limit to the number of encryption keys you can create on the Core.

Adding an encryption key

To add an encryption key:

1. Navigate to the Core Console.
2. Click **Configuration** → **Security**.
The **Encryption Keys** page appears.
3. Click **Actions**, and then click **Add Encryption Key**.
The **Create Encryption Key** dialog box displays.
4. In the **Create Encryption Key** dialog box, enter the details for the key described as follows.

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

5. Click **OK**.



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

Editing an encryption key

To edit an encryption key:

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

Changing an encryption key passphrase

To change an encryption key passphrase:

1. Navigate to the Core Console.
2. Click **Configuration** → **Security**.
The **Encryption Keys** page appears.
3. Select the encryption key you want to modify and click **Change Passphrase**.
The **Change Passphrase** dialog box appears.
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**.

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

Importing an encryption key

To import an encryption key:

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

Exporting an encryption key

To export an encryption key:

1. Navigate to the Core Console.
2. Click **Configuration** → **Security**.
3. Click > next to the name of the encryption key that you want to export, and then click **Export**.
The **Export Key** dialog box appears.
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

To remove an encryption key:

1. Navigate to Core Console.
2. Click **Configuration** → **Security**.
3. Click > next to the name of the encryption key that you want to remove, and then click **Remove**.
The **Remove Key** dialog box appears.
4. In the **Remove Key** dialog box, click **OK** to remove the encryption key.

 **NOTE:** Removing an encryption key does decrypt the data.

Managing cloud accounts

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

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

- [Removing a Cloud Account](#)

Adding a cloud account

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

To add a cloud account:

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

Table 1. Adding a cloud account

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

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

6. Click **Add**.

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

Editing a cloud account

Perform the following steps to edit a cloud account:

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



NOTE: You cannot edit the cloud type.

Configuring cloud account settings

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

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

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

Removing a Cloud account

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

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



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

Understanding replication

About protecting workstations and servers

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



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

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

About replication

Replication is the process of copying recovery points and transmitting them to a secondary location for disaster recovery. The process requires a paired source-target relationship between two cores. The source core copies the recovery points of the protected machines 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 an 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 does not prevent a recovery.
- **Replication to an Off-site Location.** The target core is 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 protected machines 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 only have access to their own data.

Possible replication configurations include:

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




Figure 7. Basic Replication Architecture Diagram

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




Figure 8. Multi-Point Replication Architecture Diagram

About seeding

Replication begins with seeding: the initial transfer of deduplicated base images and incremental snapshots of the protected machines, which can add up to hundreds or thousands of gigabytes of data. Initial replication can be seeded to the target core using external media to transfer the initial data to the target core. This is typically useful for large sets of data or sites with slow links.

NOTE: 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 protected machines.

NOTE: While replication of incremental snapshots can occur between the source and target cores before seeding is complete, the replicated snapshots transmitted from the source to the target remains “orphaned” until the initial data is consumed, and they are combined with the replicated base images.

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 the case of a severe outage in which your source core and protected machines fail, your DL Appliance supports failover and failback in replicated environments. Failover refers to switching to a redundant or standby target Core upon system failure or abnormal termination of a source core and associated protected machines. 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 a protected machine and core back to their original states (before failure). The primary goal of failback is to restore the protected machine (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.

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.

 **NOTE:** 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 must 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 1 GB 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 protected machines It is important to consider the number of protected machines per source core and how many you plan to replicate to the target. AppAssure 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 following table shows examples of the necessary bandwidth per Gigabyte for a reasonable change rate

 **NOTE:** For optimum results, adhere to the recommendations listed in the following table.

Maximum change rate for wan connection types

Table 2. Maximum change rate for wan connection types

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	838 GB per hour

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

Roadmap for performing replication

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

Performing replication in AppAssure involves performing the following operations:

- Configure self-managed replication. For more information on replicating to a self-managed target core, see [Replicating To A Self-Managed Core](#).
- Configure third-party replication. For more information on replicating to a third-party target core, see [Replicating To A Core Managed By A Third Party](#).
- Replicate a new protected machine attached to the source core. For more information on replicating a protected machine, see [Replicating A New Protected Machine](#).
- Replicate an existing protected machine. For more information on configuring an agent for replication, see [Replicating Agent Data On A Machine](#).

- Set replication priority for an agent. For more information on prioritizing the replication of agents, see [Setting Replication Priority For An Agent](#).
- Monitor replication as needed. For more information on monitoring replication, see [Monitoring Replication](#).
- Manage replication settings as needed. For more information on managing replication settings, see [Managing Replication Settings](#).
- Recover replicated data in the event of disaster or data loss. For more information on recovering replicated data, see [Recovering Replicated Data](#).

Replicating to a self-managed core

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

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

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

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

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

Text Box	Description
Host Name	Enter the host name or IP address of the Core machine to which you are replicating.
Port	Enter the port number on which the AppAssure Core communicates with the machine. The default port number is 8006.
User Name	Enter the user name for accessing the machine. For example, Administrator .
Password	Enter the password for accessing the machine.

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

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

8. If you plan to perform the seeding process for the transfer of the base data, complete the following steps:



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

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

Text Box	Description
Maximum Size	Large archives of data can be divided into multiple segments. Select the maximum size of the segment you want to reserve for creating the seed drive by doing one of the following: • Select Entire Target to reserve all available space in the path provided on the Seed Drive Location page for future use (for example, if the location is D:\work\archive, all of the available space on the D: drive is reserved if required for copying the seed drive, but is not reserved immediately after starting the copying process). • Select the blank text box, enter an amount, and then select a unit of measurement from the drop-down list to customize the maximum space you want to reserve.
Customer ID (optional)	Optionally, enter the customer ID that was assigned to you by the service provider.
Recycle action	In the event that 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 fails. • Replace this core — Overwrites any pre-existing data pertaining to this core but leave the data for other cores intact. • Erase completely — Clears all data from the directory before writing the seed drive.
Comment	Enter a comment or description of the archive.

Text Box	Description
Add all Agents to Seed Drive	Select the agents you want to replicate using the seed drive.
Build RP chains (fix orphans)	Select this option to replicate the entire recovery point chain to the seed drive. This option is selected by default. Typical seeding in AppAssure replicates only the latest recovery point to the seed drive, which reduces the amount of time and space required for creating the seed drive. Opting to build recovery point (RP) chains to the seed drive requires enough space on the seed drive to store the latest recovery points from the specified agent or agents, and may take additional time to complete the task.
Use compatible format	Select this option to create the seed drive in a format that is compatible with both new and older versions of the AppAssure Core.

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

11. Click **Finish**.

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

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

Consuming the seed drive on a target core

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

To consume the seed drive on a target core:

1. If the seed drive was saved to a portable storage device such as a USB drive, connect the drive to the target core.
2. From the Core Console on the target core, select the **Replication** tab.
3. Under **Incoming Replication**, select the correct source core by using the drop-down menu, and then click **Consume**.
The Consume window appears.
4. For **Location type**, select one of the following options from the drop-down list:
 - Local
 - Network
 - Cloud
5. Enter the following information as needed:

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

Text Box	Description
Account	Select an account from the drop-down list. To select a cloud account, you must first have added it in the Core Console.
Container	Select a container associated with your account from the drop-down menu.
Folder Name	Enter the name of the folder in which the archived data is saved; for example, - Archive-[DATE CREATED]- [TIME CREATED]

6. Click **Check File**.

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

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



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

Abandoning an outstanding seed drive

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



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

To abandon an outstanding seed drive:

1. From the Core Console on the source core, select the **Replication** tab.

2. Click **Outstanding Seed Drive (#)**.

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

3. Click the drop-down menu for the drive that you want to abandon, then select **Abandon**.

The **Outstanding Seed Drive** window appears.

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

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

Replicating to a core managed by a third party

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



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

Configuring replication to a target core managed by a third party

 **NOTE:** This configuration applies to hosted and cloud replication. If you are configuring AppAssure for multipoint to point replication, you must perform this task on all source cores.

To configure replication for a core managed by a third party:

1. Navigate to the Core Console, and 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 described as follows:

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 the protected machines to replicate.
 - b. Select a repository for each protected machine.
 - 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 displays.

8. In the **Copy to Seed Drive** dialog box, enter the information for the seed drive as 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. To designate a portion of the drive:

Text Box	Description - Enter the desired amount of space in the text box. - 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 fails. Replace this core — Overwrites any pre-existing data pertaining to this core but leave the data for other cores intact. Erase completely — Clears all data from the directory before writing the seed drive.
Comment	Enter a comment or description of the archive.
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.

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

Reviewing a replication request

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.

To review a replication request on a third-party target core:

- Open the Core Console on the target core and select the **Replication** tab.
- Click **Pending Requests (#)**.
The **Pending Replication Requests** section appears.
- Next to the request that you want to review, select **Review** from the drop-down menu.
The **Review Replication Request** window appears.



NOTE: The request completed by the customer determines the information that appears in the **Source Core Identity** section.

- On the Review Replication Request window, do one of the following:
 - To reject the request, click **Deny**.
 - To approve the request:
 - Select **Replace an existing replicated Core**, and then select a core from the drop-down list.
 - Select **Create a new source Core**. Verify the **Core Name**, customer **Email Address**, and **Customer ID**, editing the information as necessary..
- Under **Agents**, select the machines to which the approval applies, and then use select the appropriate repository for each machine by using the drop-down list.

3. Optionally, enter any notes that you want to display in the **Comment** box.
4. Click **Send Response**.

Replication is accepted.

Ignoring a replication request

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

To ignore a replication request:

1. From the Core Console on the target core, select the **Replication** tab.
2. On the Replication tab, click **Pending Requests (#)**.
The **Pending Replication Requests** section appears.
3. Next to the request that you want to ignore, select **Ignore** by using the drop-down menu.
The target core sends a notification to the source core that the request was ignored.

Monitoring replication

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

To monitor replication:

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

Table 3. Monitoring replication

Section	Description	Available Actions
Pending Replication Requests	Lists your customer ID, email address, and host name when a replication request is submitted to a third-party service provider. It is listed here until the MSP accepts the request.	In the drop-down menu, click Ignore to ignore or reject the request.
Outstanding Seed Drives	Lists seed drives that have been written but not yet consumed by the target core. It includes the remote core name, date on which it was created, and the date range.	In the drop-down menu, click Abandon to abandon or cancel the seed process.
Outgoing Replication	Lists all target cores to which the source core is replicating. It includes the remote core name, the state of existence, the number of protected machines being replicated, and the	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.

Section	Description	Available Actions
	progress of a replication transmission.	• Change Settings — Lists the display name and lets you edit the host and port for the target core. • Add Agents — Lets you select a host from a drop-down list, select protected machines for replication, and create a seed drive for the new protected machine'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.

Managing replication settings

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

To manage replication settings:

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

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

4. Click **Save**.

Removing replication

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

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

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

Removing a protected machine from replication on the source Core

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

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

Removing a protected machine on the target Core

To remove a protected machine on the target core:

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

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

Removing a target Core from replication

To remove a target core from replication:

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

Removing a source Core from replication

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

To remove a source core from replication:

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

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

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

Recovering replicated data

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

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

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

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

Roadmap for failover and failback

When you encounter a disaster situation in which your source core and associated protected machine have failed, you can enable failover in AppAssure to switch protection to your identical failover (target) core and launch a new (replicated) agent identical to the failed agent. After 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, failover and failback involve the following procedures.

- Setting up your environment for failover.
- Perform failover for the target core and associated agent.
- Restore a source core by performing failback.

Setting up an environment for failover

Setting up your environment for failover requires that you have a source and target 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 a Core for the source and install a Core for the target.
2. Install an AppAssure Agent to be protected by the source core.
3. Create one repository on the source core and one repository on the target core.
For more information, see [Creating A Repository](#).
4. Add the agent for protection under the source core.
For more information, see [Protecting A Machine](#).
5. Set up replication from the source to target core and replicate the protected agent with all recovery points.

Follow the steps in the [Replicating To A Self-Managed Core](#) 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 protected machines have failed, you can enable failover to switch protection to your identical failover (target) core. The target core becomes the only core protecting the data in your environment, and you then launch a new agent to temporarily replace the failed agent.

To perform failover on the target core:

1. Navigate to the Core Console on the target core, and click the **Replication** tab.
2. Under **Incoming Replication**, select the source core, and then expand the details under the individual agent.
3. On the **Actions** menu for that core, click **Failover**.
The 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.
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 is displayed 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](#).
11. You can now proceed with performing failback.
For more information, see [Performing Failback](#).

Performing failback

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

To perform failback:

1. Navigate to the Core Console on the target core, and click the **Replication** tab.
2. Under **Incoming Replication**, select the failover agent and expand the details.
3. On the **Actions** menu, click **Failback**.
The **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.
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.
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.

 **NOTE:** When you launch the restore, you must use the recovery points that were imported from the target core to the agent on the virtual machine.
16. Wait for the BMR reboot and for the agent service to restart, and then view and record the network connection details of the machine.
17. Navigate to the Core Console on the source core, and, on the **Machines** tab, modify the machine protection settings to add the new network connection details.
18. Navigate to the Core Console on the target core, and delete the agent from the **Replication** tab.
19. In the Core Console of the source core, set up replication again between the source and target by clicking the **Replication** tab, and then adding the target core for replication.

Managing events

Managing core events assists with the monitoring of the health and usage of the 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, e-mail SMTP settings, repetition reduction, and event retention. The Notification Groups option allows you to manage notification groups, from which you can:

- Specify an event for which you want to generate an alert for the following:
 - Clusters
 - Attachability
 - Jobs
 - Licensing
 - Log Truncation
 - Archive
 - Core Service
 - Export
 - Protection
 - Replication
 - Rollback
 - SMTP Server Settings
 - Enabled Trace logs
 - Cloud Configuration
- Specify the type of alert (error, warning, and informational).
- Specify to whom and where the alerts are sent. Options include:
 - Email Address
 - Windows Events Logs

- Syslog Server
- Specify a time threshold for repetition.
- Specify the retention period for all events.

Configuring notification groups

To configure notification groups:

1. From the Core, select the **Configuration** tab.
2. From the **Manage** option, click **Events**.
3. Click **Add Group**.

The **Add Notification Group** dialog box opens and displays three panels:

- **General**
- **Enable Events**
- **Notification Options**

4. In the **General** panel, enter basic information for the notification group described as follows:

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

5. In the **Enable Events** panel, select the conditions for which event logs (alerts) to create and report.

You can elect to create alerts for:

- **All Events**
- **Appliance Events**
- **Boot CD**
- **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** panel, specify how to handle the notification process.
The notification options are:

Text Box	Description
Notify by e-mail	Designate the recipients of the email notification. You can 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 must be reported through the Windows Event Log.
Notify by sys logd	Select this option if you want alerts to be reported through sys logd. Specify the details for the sys logd in the following text boxes: • Hostname: • Port:1

7. Click **OK**.

Configuring an email server and email notification template

If you want to receive email notifications about events, configure an email server and an email notification template.

 **NOTE:** 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 For System Events' in *Dell DL4000 Appliance User's Guide*.

To configure an email server and email notification template:

1. From the Core, select 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.
4. Select **Enable Email Notifications**, and then enter details for the email server described as follows:

Text Box	Description
SMTP Server	Enter the name of the email server to be used by the email notification template. The naming convention includes the host name, domain, and suffix; for example, smtp.gmail.com .
Port	Enter a port number. It is used to identify the port for the email server; for example, the port 587 for Gmail.

Text Box	Description
	The default is 25.
Timeout (seconds)	To specify how long to try a connection before timing out, enter an integer value. It is used to establish the time in seconds when trying to connect to the email server before a time-out occurs. The default is 30 seconds.
TLS	Select this option if the mail server uses a secure connection such as Transport Layer Security (TLS) or Secure Sockets Layer (SSL).
Username	Enter a user name for the email server.
Password	Enter a password for accessing the email server.
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 .
Email Subject	Enter a subject for the email template. It is used to define the subject of the email notification template; for example, <hostname> - <level> <name>.
Email	Enter information for the body of the template that describes the event, when it occurred, and the severity.

5. Click **Send Test Email** and review the results.
6. After you are satisfied with the results of the tests, click **OK**.

Configuring repetition reduction

To configure repetition reduction:

1. From the Core, click the **Configuration** tab.
2. From the **Manage** option, click **Events**.
3. From the **Repetition Reduction** area, click **Change**.
The Repetition Reduction dialog box appears.
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

To configure event retention:

1. From the Core, click the **Configuration** tab.
2. From the **Manage** option, click **Events**.
3. Under **Database Connection Settings**, click **change**.
The **Database Connection Settings** dialog box appears.
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, you could select 30 days (default).

5. Click **Save**.

Managing recovery

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

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

About system information

AppAssure lets you view information about the 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.

Viewing system information

To view system information:

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

Downloading installers

You can download installers from the 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 And Installing The Agent Installer](#). For more information about deploying the Agent Installer, see the *Dell DL4000 Appliance Deployment Guide* available at Dell.com/support/home. For access to the Local Mount Utility Installer, see [About The Local Mount Utility](#) and for more information about the Local Mount Utility, see [Downloading And Installing The Local Mount Utility](#).

About the agent installer

The Agent installer is used to install the AppAssure Agent application on machines that are intended to be protected by the 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 Core.



NOTE: The downloading of the Core is performed from the License Portal. To download the Core installer, visit <https://licenseportal.com>.

Downloading and installing the agent installer

You can download and deploy the Agent Installer on any machine that is protected by the Core.

To download and install the agent installer:

1. Download the Agent installer file from the License Portal or from the Core.
For example: **Agent-X64-5.3.x.xxxx.exe**
2. Click **Save File**.
For more information about installing the agents, see the *Dell DL4000 Appliance Deployment Guide* available at Dell.com/support/home.

About the local mount utility

The Local Mount Utility (LMU) is a downloadable application that lets you mount a recovery point on a remote 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 is displayed on the machine's desktop.

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

Downloading and installing the local mount utility

To download and install the Local Mount Utility:

1. From the machine on which you want to install the LMU, access the Core Console by entering the console URL into your browser and logging on with your user name and password.
2. From the 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 on the **LocalMountUtility-Web** executable and click **Open**.
Depending on your machine's configuration, the **User Account Control** window may be displayed.
7. If the **User Account Control** window is displayed, 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.

To add a core to the Local Mount Utility:

1. From the machine on which the LMU is installed, launch LMU by double-clicking the desktop icon.
2. If the **User Account Control** window displays, 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** window, enter the requested credentials described as follows:

Text Box	Description
Host name	The name of the Core from which you want to mount recovery points.  NOTE: If installing the LMU on a core, 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.

Exploring a mounted recovery point by using the local mount utility

 **NOTE:** 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 by LMU is installed, launch 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.

Mounting a recovery point by using the local mount utility

Before mounting a recovery point, the LMU must connect to the Core on which the recovery point is stored. As described in [Adding A Core To The Local Mount Utility](#), 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.

To mount a recovery point by using the Local Mount Utility:

1. From the machine on which the LMU is installed, launch LMU by double-clicking the desktop icon.
2. From the main **AppAssure Local Mount Utility** window, expand the desired core in the navigation tree to reveal the protected agents.
3. From the navigation tree, select the desired agent.
The recovery points display in the main frame.
4. Expand the recovery point that you want to mount to reveal individual disk volumes or databases.
5. Right-click the recovery point that you want to mount and select one of the following options:
 - Mount
 - Mount writable
 - Mount with previous writes
 - Advanced mount
6. From the **Advanced Mount** window, complete the options described as follows:

Text Box	Description
Mount point path	To select a path for the recovery points other than the default mount point path, click the Browse button.
Mount type	Select one of the following options: • Mount read-only • Mount writable • Mount read-only with previous writes

7. Click **Mount**.

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

 **NOTE:** Selecting a recovery point that is already mounted causes the **Mounting** dialog box to prompt you to dismount the recovery point.

Dismounting a recovery point by using the local mount utility

To dismount a recovery point using the Local Mount Utility:

1. From the machine on which the LMU is installed, launch 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. Select one of the options described in the table below to dismount recovery points.

Option	Description
Dismount	Dismounts only the adjacent recovery point. a. Click Dismount beside the chosen recovery point. b. Close the window.
Dismount all	Dismounts all mounted recovery points. a. Click Dismount all. b. In the Dismount All window, click Yes to confirm. c. Close the window.

About the local mount utility tray menu

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

Browse Recovery Points	Opens the LMU main screen.
Active Mounts	Opens the Active Mounts screen.
Options	Opens the Options screen, where you can change the Default Mount Point Directory , Default Core Credentials , and Language for the LMU user interface.
About	Opens the splash screen of licensing information.
Exit	Closes the application.

 **NOTE:** Using the X in the upper corner of the main screen minimizes the application to the tray.

Using Core and agent options

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

- Localhost Options
- Remote Core Options
- Agent Options

Accessing localhost options

To access Localhost options, right-click the 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

To access remote core options, right-click the Core or agent and then select one of the remote core options described as follows:

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

To access agent options, right-click the 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](#).

Archiving to a cloud

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

To export an archive to a cloud:

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

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 noncompliance data. The archive feature in AppAssure is used to support the extended retention for compliance and noncompliance data. It is also used to seed replication data to a remote replica core.

Creating an archive

To create an archive

1. In the Core Console, click the **Configuration** tab.
2. From the **Manage** option, click **Archive**.
The **Create Archive** dialog box appears.
3. In the **Create Archive** dialog box, enter the details for the archive described as follows:

Text Box	Description
Date range	To specify the date range, select the to and from dates.
Archive password	Enter a password for the archive, which is used to establish log in credentials to secure the archive.
Confirm	Re-enter the password to secure the archive, which is used to provide validation of the information that you entered in the Archive Password text box.
Output Location	Enter the location for the output, 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, enter a user name and password for connecting to the share.
User name	Enter a user name, which is used to establish logon credentials for the network share.
Password	Enter a password for the network path, which 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 • Specific amount in MB or GB
Recycle action	Select the appropriate recycle action.
Comment	Enter any additional information that is necessary to capture for the archive.

4. Click **Archive**.

Setting a scheduled archive

The Scheduled Archive feature lets you set a time when an archive of a selected machine will be automatically created and saved to the specified location. This accommodates situations where you would want frequent archives of a machine to be saved, without the inconvenience of needing to create the archives manually. Complete the steps in the following procedure to schedule automatic archiving. To set a scheduled archive:

1. In the Core Console, click the **Tools** tab.
2. From the **Archive** option, click **Scheduled**.
3. On the Scheduled Archive page, click **Add**.
The **Add Archive Wizard** dialog box appears.

4. On the **Location** page of the **Add Archive Wizard**, select one of the following options from the **Location Type** drop-down list:
 - Local: Output location – Enter the location for the output. It defines the location path where you want the archive to reside.
 - Network
 - Output location: Enter the location for the output. It defines the location path where you want the archive to reside.
 - User Name: Enter a user name. It establishes logon credentials for the network share.
 - Password: Enter a password for the network path. It establishes logon credentials for the network share.
 - Cloud
 - Account : Select an account from the drop-down list. To select a cloud account, you must first have added it in the Core Console.
 - Container: Select a container associated with your account from the drop-down menu.
 - Folder Name: Enter a name for the folder in which the archived data is to be saved. The default name is AppAssure-5-Archive-[DATE CREATED]-[TIME CREATED]
5. Click **Next**.
6. On the **Machines** page of the wizard, select which protected machines contain the recovery points you want to archive.
7. Click **Next**.
8. On the **Options** page, select one of the following Recycle Actions from the drop-down list:
 - **Replace this Core**: Overwrites any existing archived data pertaining to this core but leaves the data for other cores intact.
 - **Erase completely**: Clears all archived data from the directory before writing the new archive.
 - **Incremental**: Lets you add recovery points to an existing archive. It compares recovery points to avoid duplicating data that already exists in the archive.
9. On the **Schedule** page, select one of the following Send data frequency options:
 - Daily: At time – Select the hour of the day you want to create a daily archive.
 - Weekly
 - At day of week: Select a day of the week on which to automatically create the archive.
 - At time: Select the hour of the day you want to create a daily archive.
 - Monthly
 - At day of months: Select the day of the month on which to automatically create the archive.
 - At time: Select the hour of the day you want to create a daily archive.
10. To pause archiving for resuming at a later time, select **Initial pause archiving**.

You may want to pause the scheduled archive if you need time to prepare the target location before archiving resumes. If you do not select this option, archiving begins at the scheduled time.
11. Click **Finish**.

Pausing or resuming scheduled archive

If you opted to initially pause archiving when you performed the Setting a Scheduled Archive procedure, you would want to resume the scheduled archive at a later time.

To pause or resume scheduled archive:

1. Navigate to the **Core Console**, and then click the **Tools** tab.
2. From the **Archive** option, click **Scheduled**.
3. On the **Scheduled Archive** page, do one of the following:
 - Select the preferred archive, and then click one of the following actions as appropriate:
 - Pause
 - Resume
 - Next to the preferred archive, click the drop-down menu, and then click one of the following actions as appropriate:
 - Pause
 - Resume

The status of the archive displays in the **Schedule** column.

Editing a scheduled archive

1. In the Core Console, click the **Tools** tab.
2. From the **Archive** option, click **Scheduled**.
3. On the Scheduled Archive page, click the drop-down menu next to the archive you want to change, and then click **Edit**.

The **Add Archive Wizard** dialog box appears.
4. On the **Location** page of the **Add Archive Wizard**, select one of the following options from the **Location Type** drop-down list:
 - Local: Output location – Enter the location for the output. It defines the location path where you want the archive to reside.
 - Network
 - Output location: Enter the location for the output. It defines the location path where you want the archive to reside.
 - User Name: Enter a user name. It establishes logon credentials for the network share.
 - Password: Enter a password for the network path. It establishes logon credentials for the network share.
 - Cloud
 - Account : Select an account from the drop-down list. To select a cloud account, you must first have added it in the Core Console.
 - Container: Select a container associated with your account from the drop-down menu.
 - Folder Name: Enter a name for the folder in which the archived data is to be saved. The default name is AppAssure-5-Archive-[DATE CREATED]-[TIME CREATED]
5. Click **Next**.
6. On the **Machines** page of the wizard, select which protected machines contain the recovery points you want to archive.
7. Click **Next**.
8. On the **Schedule** page, select one of the following Send data frequency options:
 - Daily: At time – Select the hour of the day you want to create a daily archive.
 - Weekly

- At day of week: Select a day of the week on which to automatically create the archive.
- At time: Select the hour of the day you want to create a daily archive.
- Monthly
 - At day of months: Select the day of the month on which to automatically create the archive.
 - At time: Select the hour of the day you want to create a daily archive.
- 9. To pause archiving for resuming at a later time, select **Initial pause archiving**.
You may want to pause the scheduled archive if you need time to prepare the target location before archiving resumes. If you do not select this option, archiving begins at the scheduled time.
- 10. Click **Finish**.

Checking an archive

You can scan an archive for structural integrity by performing an archive check. This check verifies the presence of all necessary files within the archive. To perform an archive check, complete the steps in the following procedure:

1. In the Core Console, click the **Tools** tab.
2. From the **Archive** option, click **Check Archive**.
The **Check Archive** dialog box appears.
3. Select one of the following options from the drop-down list:
 - Local: Output location – Enter the location for the output. It defines the location path where you want the archive to reside.
 - Network
 - Output location: Enter the location for the output. It defines the location path where you want the archive to reside.
 - User Name: Enter a user name. It establishes logon credentials for the network share.
 - Password: Enter a password for the network path. It establishes logon credentials for the network share.
 - Cloud
 - Account : Select an account from the drop-down list. To select a cloud account, you must first have added it in the Core Console.
 - Container: Select a container associated with your account from the drop-down menu.
 - Folder Name: Enter a name for the folder in which the archived data is to be saved. The default name is AppAssure-5-Archive-[DATE CREATED]-[TIME CREATED]
4. To also perform a structure integrity check, select **Structure integrity**.
5. Click **Check File**.

Importing an archive

To import an archive:

1. In the Core Console, select the **Configuration** tab.
2. From the **Manage** option, click **Archive**, and then click **Import**.
The **Import Archive** dialog box appears.
3. In the **Import Archive** dialog box, enter the details for importing the archive described as follows:

- | Text Box | Description |
|-----------------------|---|
| Input Location | Select the location for importing the archive. |
| User name | To establish access to secure the archive, enter the logon credentials. |
| Password | Enter a password for accessing the archive. |
- Click **Check File** to validate the existence of the archive to import.
The **Restore** dialog box appears.
 - In the **Restore** dialog box, verify the name of the source core.
 - Select the agents to import from the archive.
 - Select the repository.
 - Click **Restore** to import the archive.

Managing SQL attachability

The SQL attachability configuration enables the 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, 2012 and 2014. 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.

 **CAUTION:** The version of SQL Server on the Core must be equal to or newer than the SQL Server version on all the agents with SQL Server installed.

Configuring SQL attachability settings

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

 **NOTE:** 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.

To configure SQL attachability settings:

- Navigate to the Core Console. click the tab.
- Click **Configuration** → **Settings**.
- In the Nightly Jobs pane, click **change**.
The **Nightly Job** dialog box appears.
- Select **Attachability Check Job** and then click **Settings**.

5. Use the drop-down menus to select the instance of SQL Server installed on the Core from the following options:

You can choose from:

- **SQL Server 2005**
- **SQL Server 2008**
- **SQL Server 2008 R2**
- **SQL Server 2012**
- **SQL Server 2014**

6. Select the credential type.

You can choose from:

- **Windows**
- **SQL**

7. Specify the credentials with administrative privileges for the Windows or SQL Server instances, described as follows:

Text Box	Description
Username	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.

8. Click **Test Connection**.



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

9. Click **Save**.

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

10. In the Nightly Jobs window, click **OK**.

Attachability checks are now schedule to occur with the nightly jobs.

Configuring nightly SQL attachability checks and log truncation

To configure nightly SQL attachability checks and log truncation:

1. In the left navigation area of the Core, select the machine for which you want to have nightly attachability checks and log truncation, and click **SQL Server Settings**.
2. Navigate to the Core Console.
3. Click **Configuration** → **Settings**.
4. In the **Nightly Jobs** section, click **Change**.
5. Select or clear the following SQL Server settings based on the needs of your organization:
 - **Attachability Check Job**
 - **Log Truncation Job (simple recovery model only)**
6. Click **OK**.

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

Managing exchange database mountability checks and log truncation

When using AppAssure 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 is recovered successfully in the event of a failure.

 **NOTE:** The mountability checks and log truncation features only apply to Microsoft Exchange 2007, 2010, and 2013. Additionally, the AppAssure 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.

To configure Exchange database mountability and log truncation:

1. In the left navigation area of the Core Console, select the machine for which you want to configure mountability checks and log truncation.
The **Summary** tab for the selected machine is displayed.
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](#).

Forcing a mountability check

To force a mountability check:

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

 **NOTE:** For instructions on how to view the status of the attachability checks, see [Viewing Events And Alerts](#).

The system performs the mountability check.

Forcing checksum checks

To force a checksum check:

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

The **Force Attachability Check** window prompts you to indicate if 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](#).

Forcing log truncation



NOTE: This option is only available for Exchange or SQL machines.

To force log truncation:

1. Navigate to the 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

After a recovery point is created on a protected SQL or Exchange server, the application is displayed a corresponding color status indicator in the **Recovery Points** table. The color that is displayed is based on the check settings for the protected machine and the success or failure of those checks, as described in the following tables.



NOTE: For more information on viewing Recovery Points, see [Viewing Recovery Points](#).

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

Recovery Status Point Colors 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. • 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.

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

Recovery Status Point Colors for Exchange Databases

Term heading	Description heading
White	Indicates that one of the following conditions exist:

Term heading	Description heading
	• An Exchange database did not exist. • 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 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.
	NOTE: Recovery points that do not have an Exchange or SQL database associated with it is displayed with a white status indicator. In situations where both an Exchange and SQL database exists for the recovery point, the most severe status indicator is displayed for the recovery point.

Managing Your Appliance

The Core Console includes an **Appliance** tab, which you can use to provision space, monitor the health of the appliance, and access management tools.

Monitoring the status of the Appliance

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

The Overall Status page also provides links to tools that drill down into the details of each subsystem, which can be helpful for troubleshooting warnings or errors. The **System Administrator** link, available for the Appliance Hardware and Storage Hardware subsystems, prompts you to log on to the System Administrator application used for managing hardware. For more information about the System Administrator application, see the *OpenManage Server Administrator User's Guide* on dell.com/support/home. The **Provisioning Status** link, available for the Storage Provisioning subsystem, opens the **Tasks** screen which displays the provisioning status of that subsystem. If storage is available for provisioning, a link to **Provision** under **Actions** displays next to the provision task.

Provisioning storage

The appliance configures available DL4000 internal storage and any attached external storage enclosures for:

- AppAssure Repositories

 **NOTE:** If fibre channel HBA is configured then the process of creating the repositories is manual. AppAssure will not create a repository automatically in the root directory. For more information, see the *Dell DL4000 Appliance Deployment Guide*.

- Virtual Standby of Protected Machines

 **NOTE:** MD1400s with 1 TB, 2 TB, 4 TB or 6 TB (for high capacity) drives connected to the H830 controller are supported. Up to four MD 1400s are supported.

 **NOTE:** The DL4000 high-capacity configuration supports either H810 PERC SAS adapter or two Fibre Channel HBAs. For more information on configuring fibre channel HBAs, see the *DL4xxx — Fibre Channel Implementation* whitepaper located at dell.com/support/home.

Before you begin provisioning storage on the disk, determine how much storage you want for standby virtual machines. You can allocate any percentage of the available capacity to host standby virtual machines. For example, if you are using Storage Resource Management (SRM), you can allocate up to 100 percent capacity on any device being provisioned to host virtual machines. Using AppAssure's Live Recovery feature, you can use these virtual machines to quickly replace any failed server that the appliance protects.

Based on a medium-sized environment that does not need standby virtual machines, you can use all of the storage to back up a significant number of agents. However, if you need more resources for standby virtual machines and back up a smaller number of agent machines, you can allocate more resources for larger VMs.


When you select the **Appliance** tab, the AppAssure Appliance software locates the available storage space for all supported controllers in the system and validates that the hardware meets the requirements.

To complete disk provisioning for all available storage:

1. In the **Appliance** tab, click **Tasks** → **Provisioning**.
The **Provisioning** screen displays estimated capacity for provisioning. This capacity is used to create a new AppAssure Repository.

 **CAUTION:** Before proceeding ensure Step 2 through Step 4 is followed in this procedure.

2. Open the **Provisioning Storage** window by clicking **Provision** in the Action column next to the storage that you want to provision.
3. In the **Optional Storage Reserve** section, select the box next to **Allocate a portion of the storage being provisioned for Standby Virtual Machines or other purposes** and indicate a percentage of storage to allocate. Otherwise, the percentage of storage indicated in the **Optional Storage Reserve** section will be taken from all of the attached disks.
4. Click **Provision**.



Provisioning selected storage

To provision selected storage:

1. In the **Appliance** tab, click **Tasks** → **Provisioning**.
The **Provisioning** screen displays estimated capacity for provisioning. This capacity is used to create a new AppAssure Repository.
2. To provision only a portion of the available space, click **Provision** under **Action** next to the storage space that you want to provision.
 - To create new repository, select **Create a new repository**, and provide a name for the repository. By default, Repository 1 appears as the repository name. You can opt to overwrite the name.
 - To add capacity to an existing repository, select **Expand the existing repository**, and then select the repository from the **Existing Repositories** list.

 **NOTE:** To add capacity, it is recommended that you expand an existing repository instead of adding a repository. Separate repositories do not utilize capacity as efficiently because deduplication cannot occur across separate repositories.
3. Under **Optional Storage Reserve**, select **Allocate a portion of the storage being provisioned for Standby Virtual Machines or other purposes**, and then specify the percentage of storage to allocate for the VMs.

4. Click **Provision**.

The disk provisioning begins and the status of the AppAssure repository creation is displayed in the **Status** area of the **Tasks** screen. The **State** displays **Provisioned**.

5. To view the details after disk provisioning completes, click > next to the status light.

The **Tasks** page expands and displays status, repository, and virtual disk details (if allocated).

Deleting space allocation for a virtual disk

Before you begin this procedure, determine which virtual disk you want to delete. From the Core Console, select the **Appliance** tab, click **Tasks**, and then expand the repository that contains the virtual disks to see the virtual disk details.

To delete a space allocation for a virtual disk:

1. From the OpenManage Server Administrator application, expand **Storage**.
2. Expand the controller that houses the virtual disk, then select **Virtual Disks**.
3. Select the virtual disk that you want to remove, and then select **Delete** from the **Tasks** drop-down menu.
4. After confirming the deletion, the space appears on the Core Console **Appliance** tab **Tasks** screen as available for provisioning.

Resolving failed tasks

AppAssure reports failed verify, provision, and recovery tasks with an event on the Core Console Home page, and also on the **Appliance** tab **Tasks** screen.

To understand how to resolve a failed task, select the **Appliance** tab and then click **Tasks**. Expand the failed task by clicking > next to **Status**, and review the error message and recommended action.

Upgrading your Appliance

To upgrade your appliance:

1. Download the **Recovery and Update Utility** from dell.com/support to the DL4000 Backup to Disk appliance.
2. Copy the utility to the appliance desktop and extract the files.
3. Double-click the **launchRUU** icon.
4. When prompted, click **Yes** to acknowledge that you are not running any of the listed processes.
5. When the **Recovery and Update Utility** screen appears, click **Start**.
6. When prompted to reboot, click **OK**.

The updated versions of the Windows Server Roles and Features, ASP .NET MVC3, LSI Provider, DL Applications, OpenManage Server Administrator and AppAssure Core Software are installed as part of the Recovery and Update Utility. In addition to these, the Recovery and Update Utility also updates the RASR content.



NOTE: As part of the AppAssure Core Software upgrade process, the Recovery and Upgrade Utility notifies you of the currently installed AppAssure version and prompts you to confirm that you want to upgrade the Core software to the version that is bundled in the utility. AppAssure Core software downgrades are not supported.

7. If prompted, reboot your system.
8. After all services and applications are installed, click **Proceed**.

The Core Console launches.

Repairing your Appliance

To repair your appliance:

1. Download the **Recovery and Update Utility** from **dell.com/support** to your Appliance.
2. Copy the utility to the appliance desktop and extract the files.
3. Double-click the **launchRUU** icon.
4. When prompted, click **Yes** to acknowledge that you are not running any of the listed processes.
5. When the Recovery and Update Utility screen displays, click **Start**.
6. When prompted to reboot, click **OK**.

The updated versions of the Windows Server Roles and Features, ASP .NET MVC3, LSI Provider, DL Applications, OpenManage Server Administrator and AppAssure Core Software are installed as part of the Recovery and Update Utility.

7. If the bundled version in the utility is the same as the installed version, the Recovery and Update Utility prompts you to confirm that you want to run a repair installation. This step can be skipped if a repair install of the AppAssure Core is not needed.
8. If the bundled version in the utility is higher than the installed version, then the Recovery and Update Utility prompts you to confirm that you want to upgrade the AppAssure Core software.



NOTE: AppAssure Core software downgrades are not supported.

9. If prompted, reboot your system.
10. After all services and applications are installed, click **Proceed**.

AppAssure Appliance Configuration Wizard will be launched if the system needs to be configured again after repair, otherwise Core Console will be launched.

Protecting workstations and servers

About protecting workstations and servers

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

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

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

Configuring machine settings

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

Viewing and modifying configuration settings

To view and modify configuration settings:

1. After you have added a protected machine, perform one of the following:
 - From the Core Console, click the **Machines** tab and then click the hyperlink for the machine that you want to modify.
 - From the **Navigation** pane, select the machine that you want to modify.
2. Click the **Configuration** tab.
The **Settings** page displays.
3. Click **Edit** 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 Core Console. By default, this is the host name of the machine. You can change the display name to something more user-friendly if needed.
Host Name	Enter a host name for the machine.
Port	Enter a port number for the machine.

Text Box	Description The Core uses the port to communicate with this machine.
Repository	Select a repository for the recovery points. Displays the repository on the 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.
Encryption Key	Edit the encryption key if necessary. Specifies whether encryption is applied to the data for every volume on the machine that is stored in the repository.

Viewing system information for a machine

The Core Console displays all the machines that are being protected by including a list of the machines as well as each machine's status.

To view system information for a machine:

1. In the Core Console, under **Protected Machines**, select the machine for which you want to view detailed system information.
2. Click the **Tools** tab for that machine.

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
- Virtual Machine Type (if applicable)

Detailed information about the volumes contained on this machine includes:

- Name
- Device ID
- File System
- Capacity (including Raw, Formatted, and Used)
- Processors
- Type of processors
- Network adapters
- IP addresses associated with this machine

Configuring notification groups for system events

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

To configure notification groups for system events:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to modify.
 - In the navigation pane, select the machine that you want to modify.

The **Summary** tab appears.

3. Click the **Configuration** tab, and then click **Events**.
The **Notification Groups** page appears.
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.
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

Text Box	Description • 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 — Specify to which email addresses to send the events in the To, CC, and 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 — 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.

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 opens where you can edit the settings.

Editing notification groups for system events

To edit notification groups for system events:

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

Text Box

Description

Description

Enter a description for the notification group.

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**
- **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** — Specify the e-mail addresses to which to send the events in the To, CC, and BCC text boxes.



NOTE: To receive e-mail, SMTP must be previously configured.

- **Notify by Windows Event log** — The Windows Event log controls the notification.
- **Notify by syslogd** — You must specify the host name and port to which to send the events.
 - **Host** — Enter the host name for the server.

Text Box	Description
	– Port — Enter a port number for communicating with the server.

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. A rollup process enforces the retention policy, and helps with aging and deleting old backups. This task is also a step in [Process Of Modifying Cluster Node Settings](#).

To customize retention policy settings:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to modify.
 - In the navigation pane, select the machine that 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, ensure that you select the Use Core default retention policy option.

The **Retention Policy** screen appears.

4. To set the customized policies, click **Use custom retention policy**.

The **Custom Retention Policy** screen appears.

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

Text Box	Description
Keep all Recovery Points for n [retention time period]	Specifies the retention period for the recovery points. Enter a number that represents 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 refined 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 that represents the retention period and then select the time period. The default is 2. You can choose from: • Days

Text Box	Description
	• Weeks • Months • Years
...and then keep one Recovery Point per day for n [retention time period]	Provides a more refined level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number that represents 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 refined level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number that represents the retention period and then select the time period. The default is 3. You can choose from: • Weeks • Months • Years
...and then keep one Recovery Point per month for n [retention time period]	Provides a more refined level of retention. It is used as a building block to further define how long recovery points are maintained. Enter a number that represents 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 that represents the retention period and then select the time period.

The Newest Recovery Point text box is displayed the most recent recovery point. The retention policy settings determine the oldest recovery point.

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 can be one year, four months, and six days old.

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

Viewing license information

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

To view license information:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to view.
 - In the navigation pane, select the machine that you want to view.
3. Click the **Configuration** tab, and then click **Licensing**.
The **Status** screen displays the details about the product licensing.

Modifying protection schedules

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

To modify protection schedules:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to modify.
 - In the navigation pane, select the machine that you want to modify.
3. Do one of the following:
 - In the **Volumes** table on the **Summary** tab for the machine, click the hyperlink for the protection schedule for the volume that you want to customize.
 - Click the **Configuration** tab, and then click **Protection Settings**. In the list of volumes, click the **Edit** icon next to the volume you want to customize.

The **Protection Schedule** dialog box appears.

4. In the **Protection Schedule** dialog box, edit the following schedule options as needed for protecting your data. The following table describes the options.

Option	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, 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 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**.

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

Modifying transfer settings

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

 **CAUTION:** Changing transfer setting could have dramatic effects on your environment. Before modifying transfer settings values, refer to the Transfer Performance Tuning Guide in the Dell AppAssure knowledge base <https://support.software.dell.com/appassure/kb>.

There are three types of transfers:

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

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

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

- Number of concurrent agent data transfers
- Number of concurrent data streams

- Amount of data change on disk
- Available network bandwidth
- Repository disk subsystem performance
- Amount of memory available for data buffering

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

To modify transfer settings:

1. In the Core Console, do one of the following:
 - Click the **Machines** tab, and then click the hyperlink for the machine that you want to modify.
 - In the navigation pane, click the machine that you want to modify.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to modify.
 - In the navigation pane, select the machine that you want to modify.
3. Click the **Configuration** tab, and then click **Transfer Settings**.
The current transfer settings appear.
4. On the **Transfer Settings** page, click **Change**.
The **Transfer Settings** dialog box appears.
5. 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 value to the same value that you select for Maximum Concurrent Streams. If you experience packet loss, set this value slightly lower. For example, if Maximum Current Streams is set at 8, set this option to 7.
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.

Text Box	Description
	 CAUTION: Do not change this option from the default setting.
Maximum Transfer Queue Depth	Specifies the number of commands that can be sent concurrently. You can adjust this option 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.
Excluded Writers	Select a writer if you want to exclude it. Since the writers that appear in the list are specific to the machine that you are configuring, you may not see all the writers listed. Some writers you may see include: • 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 is not performed in that time, the operation is repeated.
Network Write Timeout	Specifies the maximum time in seconds to wait for a write connection. If the network write is not performed in that time, the operation is repeated.

6. Click **OK**.

Restarting a service

To restart a service:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:

- Click the hyperlink for the machine that you want to restart.
 - In the **Navigation** pane, select the machine that 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, view the logs to troubleshoot.

To view machine logs:

1. In the Core Console, click the **Machines** tab.
2. From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine containing the logs that you want to view.
 - In the **Navigation** pane, select the machine containing the logs that you want to view.
3. Click the **Tools** tab, and then click **Diagnostics**.
4. Click the **View Log** link.

Protecting a machine

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

 **NOTE:** The machine must have the 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](#).

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

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

To protect a machine:

1. If you did not do so after installing the Agent software, reboot the machine on which the Agent software is installed.
2. From the Core Console on the core machine, 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 appears.

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 Core communicates with the agent on the machine. The default port number is 8006.
Username	The user name used to connect to this machine; for example, administrator.

Text Box	Description
----------	-------------

Password	The password used to connect to this machine.
-----------------	---

4. Click **Connect** to connect to this machine.



NOTE: If the Agent software is not yet installed on the machine that you designated, follow the procedure [Deploying The Agent Software When Protecting An Agent](#). Restart the agent machine after deploying the Agent software, and then resume with the next step.

5. In the **Protect** dialog box, edit the settings as needed as described in the following table.

Field	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 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 Core in which to store the data from this machine.
-------------------	---

Encryption Key	Specify whether encryption is 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 Core Console.

Initially Pause Protection	After you add a machine for protection, AppAssure automatically begins taking a base snapshot of data. You can select this check box to pause protection initially. You then must force a snapshot manually when you are ready to start protecting your data. For more information about manually forcing a snapshot, see Forcing A Snapshot .
-----------------------------------	--

Volume Groups	Under Volume Groups, you can define which volumes that 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](#).

6. Click **Protect**.

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

 **CAUTION:** If you protected a Linux machine, you must not unmount a protected volume manually. 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 must be in a form similar to this example: `/dev/sda1`.

Deploying the agent software when protecting an agent

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

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

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

1. From the **Protect Machine** → **Connect** dialog box, after entering the appropriate connection settings, click **Connect**.
The **Deploy Agent** dialog box is displayed.
2. Click **Yes** to deploy the Agent software remotely to the machine.
The **Deploy Agent** dialog box is displayed.
3. Enter logon and protection settings as follows:
 - **Host name** — Specifies the host name or IP address of the machine that you want to protect.
 - **Port** — Specifies the port number on which the Core communications with the Agent on the machine. The default value is 8006.
 - **User name** — Specifies the user name used to connect to this machine; for example, administrator.
 - **Password** — Specifies the password used to connect to this machine.
 - **Display name** — Specifies a name for the machine which appears on the Core Console. The display name could be the same value as the host name.
 - **Protect machine after install** — Selecting this option enables AppAssure to take a base snapshot of the data after you add the machine for protection. This option is selected by default. If you deselect this option, then you must force a snapshot manually when you are ready to start data protection. For more information about manually forcing a snapshot, see topic 'Forcing A Snapshot' in *Dell DL4000 Appliance User's Guide*.
 - **Repository** — Select the repository in which to store data from this agent.
 **NOTE:** You can store data from multiple agents in a single repository.
 - **Encryption Key** — Specifies whether encryption should be applied to the data for every volume on this machine to be stored in the repository.
 **NOTE:** You define encryption settings for a repository under the **Configuration** tab in the Core Console.
4. Click **Deploy**.
The **Deploy Agent** dialog box closes. There may be a delay before you see the selected agent appear in the list of protected machines.

Creating custom schedules for volumes

To create custom schedules for volumes:

1. In the **Protect Machine** dialog box (for information about accessing this dialog box, see [Protecting A Machine](#), under **Volume Groups**, select a volume for protection, and then click **Edit**.

The **Protection Schedule** dialog box appears.

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

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 that 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 must configure additional settings in the Core Console.

To modify Exchange server settings:

1. After you add the Exchange Server machine for protection, select the machine in the Core Console's **Navigation** pane.

The **Summary** tab appears 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.
 - 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 the new credentials and then click **OK**.

Modifying SQL server settings

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

To modify SQL server settings:

1. After you have added the SQL Server machine for protection, select the machine in the Core Console's **Navigation** pane.
The **Summary** tab appears 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 the new credentials, and then click **OK**.

Deploying an agent (push install)

AppAssure requires microsoft.net for the agent installation. Microsoft.net must be installed on any client machine prior to installing the agent either manually or by a push installation process.

AppAssure lets you deploy the AppAssure 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](#).

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

To deploy an agent:

1. From the Core Console, 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 machine that you want to deploy.
Username	Enter the user name to connect to this machine (for example, administrator).
Password	Enter the password to connect to this machine.

- | Text Box | Description |
|---------------------------------------|--|
| Automatic reboot after install | Select to specify whether the Core starts upon the completion of the deployment and installation of the AppAssure Agent Installer. |
- Click **Verify** to validate the credentials you entered.
The **Deploy Agent** dialog box displays a message indicating that validation is being performed.
 - Click **Abort** if you want to cancel the verification process.
After the verification process completes, a message indicating that verification has been completed appears.
 - Click **Deploy**.
A message indicating that the deployment has started appears. You can view the progress in the **Events** tab.
 - Click **Show details** to view more information about the status of the agent deployment.
 - Click **OK**.

Replicating a new agent

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

To replicate a new agent:

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

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

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

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



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

Text Box	Description
Remote Core	Specify the target core to which you want to replicate the agent.
Remote Repository	The name of the desired repository on the target core in which to store the replicated data from this machine.
Pause	Select this check box if you want to pause replication; for example, to pause it until after AppAssure takes a base image of the new agent.
Schedule	Select one of the following options: - Protect all volumes with default schedule - Protect specific volumes with custom schedule  NOTE: The default schedule is every 15 minutes.
Initially pause protection	Select this check box if you want to pause protection; for example, to prevent AppAssure from taking the base image until after peak usage hours.

- 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

- Navigate to the Core Console and then click the **Machines** tab.
- From the **Machines** tab, perform one of the following:
 - Click the hyperlink for the machine that you want to remove.
 - Or, in the navigation pane, select the machine that you want to remove.
- In the **Actions** drop-down menu, click **Remove Machines**, and then select one of the option described in the following table.

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.

Replicating agent data on a machine

Replication is the relationship between the target and source cores in the same site, or across two sites with slow link on a per agent basis. When replication is set up between two cores, the source core asynchronously transmits the incremental snapshot data of select agents to the target or source core.

Outbound replication can be configured to a Managed Service Provider providing off-site backup and disaster recovery service or to a self-managed core. To replicate agent data on a machine:

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

The **Enable Replications** dialog box appears.

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

Setting replication priority for an agent

To set replication priority for an agent:

1. From the Core Console, select the protected machine for which you want to set replication priority, and click the **Configuration** tab.
2. Click **Select Transfer Settings**, and then use the **Priority** drop-down list to select one of the following options:
 - **Default**
 - **Highest**
 - **Lowest**
 - **1**
 - **2**
 - **3**
 - **4**

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

3. Click **OK**.

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 includes exports, replications, and so on.

To cancel operations on a machine:

1. From the 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 as follows:

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

Viewing machine status and other details

To view machine status and other details:

1. In the navigation pane of the Core Console, do one of the following:
 - Select the **Machines** tab, and then click the hyperlink for the machine that you want to view.
 - In the navigation pane, click the machine that you want to view.

The **Summary** tab is displayed.

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

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

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

- Total size
- Used Space
- Free space

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

- Name
- 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 Agent software simultaneously to multiple Windows machines.

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

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



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

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

Deploying to multiple machines

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

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

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



NOTE: The machines that you are deploying must have access to the Internet to download and install bits as AppAssure uses the Web version of the AppAssure Agent Installer to deploy the installation components. If access to the Internet is not available, you can push the AppAssure Agent installation program from Core machine. For information about pushing the Agent installation from the Core machine, see [Pushing The Agent Installation Program From The Core Machine](#). You can download core and agent updates from the License Portal.

Pushing the agent installation program from the core machine

If the servers being deployed do not have internet access, then you can push the actual agent installation file from the Core machine. Your appliance includes the agent installation program file.



NOTE: Download Core and agent upgrades from the License Portal.

To push the agent installation program from the Core machine:

1. From the Core machine, copy the agent installation file **Agent-X64-5.x.x.xxxx.exe** to the **C:\Program Files\apprecovery\core\installers** directory.
2. From the Core Console, select the **Configuration** tab, and then click **Settings**.
3. In the **Deploy Settings** section, edit the **Agent Installer Name**.

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 the agent to multiple machines on an Active Directory domain:

1. From the 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 described in the following table:

Text Box	Description
Domain	The host name or IP address of the Active Directory domain.
User name	The username 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 Agent, and then click **Add**.

The machines that you added appear on the **Deploy Agent on 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 Core communicates with the agent on the machine.
User name	Automatically provided from step 3.
Password	Enter the password for the machine.
Automatic reboot after install	Specify whether you want to automatically reboot the machine after deployment.



NOTE: This option is mandatory if you want to automatically protect the machine after deployment by checking the **Protect Machine After Install** box.

Protect Machine After Install	Specify whether you want to automatically protect the machine after deployment. This allows you to skip Protecting Multiple Machines .
--------------------------------------	---

Text Box	Description
Repository	Use the drop-down list to select the repository on the Core where the data from the machines should be stored. The repository that you select is used for all of the machines that are being protected.  NOTE: This option is only available when you select Protect machine after install .
Encryption Key	(Optional) 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 of the machines that are being protected.  NOTE: This option is only available when you select Protect machine after install

b. Click **Save**.

- To verify that AppAssure can connect to each machine successfully, select each machine on 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:

Text Box	Description
Green icon	AppAssure is able to connect to the machine and it is ready to be deployed.
Yellow icon	AppAssure is able to connect to the machine; however, the agent is already paired with a core machine.
Red icon	AppAssure 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.

- After the machines are successfully verified, select each machine to which you want to deploy the AppAssure 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 or ESXi virtual host

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

 **NOTE:** All virtual machines must have VM Tools installed; otherwise, AppAssure cannot detect the host name of the virtual machine to which to deploy. In lieu of the host name, AppAssure 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/ESXi virtual host:

- From the Core Console, click the **Tools** tab, and then click **Bulk Deploy**.
- On the **Deploy Agent on Machines** window, click **vCenter/ESXi**.
- In the **Connect to VMware vCenter Server/ESXi** dialog box, enter the host information and logon credentials as follows and click **OK**.

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

4. On the **Add Machines from VMware vCenter Server/ESXi** dialog box, check the box next to the machines to which you want to deploy the AppAssure Agent, and then click **Add**.
5. On the **Deploy Agent on Machines** window, you can view the machines that you added. If you want to select a repository, encryption key, or other settings for a machine, select the check box next to the machine and click **Edit Settings**.

For details on each setting, see [Deploying To Machines On An Active Directory Domain](#).

6. Verify that AppAssure can connect to each machine successfully. Select each machine on the **Deploy Agent on Machines** window, and then click **Verify**.
7. The **Deploy Agent on Machines** window shows an icon next to each machine that reflects its readiness for deployment, as follows:

Text Box	Description
Green icon	AppAssure is able to connect to the machine and it is ready to be deployed.
Yellow icon	AppAssure is able to connect to the machine; however, the agent is already paired with a core machine.
Red icon	AppAssure 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 and click **Deploy**.
9. If you chose the **Protect machine after install** option, after deployment is successful, the machines automatically rebooted and protection is enabled.

Deploying to machines on any other host

To deploy to multiple machines on any other host:

1. From the 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 specify multiple machines by using the **Add Machine** dialog box; this allows you 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](#). 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, select the **Protect Machine after Install** check box. If you select the check box, the machine reboots automatically before enabling protection.

- Click **Manually** to specify multiple machines in a list; each line represents a machine to which to deploy. In the **Add Machines Manually** dialog box, enter the IP address or name for the machine, the user name, password separated by a double-colon delimiter, and port as follows:

```
hostname::username::password::port
For example:
10.255.255.255::administrator::&11@yYz90z::8006
abc-host-00-1::administrator::99!zU$o83r::168
```

3. On the **Deploy Agent on Machines** window, you can see the machines that you added. If you want to select a repository, encryption key, or other settings for a machine, select the check box next to the machine and click **Edit Settings**.

For details on each setting, see [Deploying To Machines On An Active Directory Domain](#).

4. Verify that AppAssure can connect to each machine successfully. Select each machine on the **Deploy Agent on 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:

Text Box	Description
Green icon	AppAssure is able to connect to the machine and it is ready to be deployed.
Yellow icon	AppAssure is able to connect to the machine; however, the agent is already paired with a core machine.
Red icon	AppAssure 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.

5. After machines are verified successfully, check the box next to each machine and click **Deploy**.
6. If you chose the **Protect machine after install** option, after deployment is successful, the machines are automatically rebooted and protection is enabled.

Monitoring the deployment of multiple machines

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

To monitor the deployment of multiple machines:

1. From the Core Console, click the **Events** tab, locate the deployment job in the list, and click the button in the **Details** column.

The **Monitor Active Task** window displays the details of the deployment.

It includes overall progress information as well as the status for each individual deployment. Details that display include:

- Start Time
- End Time
- Elapsed Time
- Time Remaining
- Progress
- Phase

2. Do one of the following:
 - Click **Open in New window** to launch a new window to view the progress of the deployment.
 - Click **Close** and the deployment tasks process in the background.

Protecting multiple machines

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



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

To protect multiple machines:

1. From the Core Console, click the **Tools** tab, and then click **Bulk Protect**.
The **Protect Machines** window appears.
2. Add the machines that you want to protect by clicking one of the following options.
For details on completing each option, see [Deploying To Multiple Machines](#).
 - Click **Active Directory** to specify machines on an Active Directory domain.
 - Click **vCenter/ESXi** to specify virtual machines on a vCenter/ESXi virtual host.
 - Click **New** to specify multiple machines by using the Add Machine dialog box.
 - Click **Manually** to specify multiple machines in a list by typing host name and credentials.
3. On the **Protect Machines** window, you can view the machines that you added. If you want to select a repository, encryption key, or other advanced settings for a machine, select the check box next to the machine and click **Edit Settings**.
4. Specify the settings as follows and click **OK**.

Text Box	Description
Username	Enter the username used to connect to this machine; for example, Administrator.
Password	Enter the secure password used to connect to this machine.
Port	Specify the port number on which the Core communicates with the agent on the machine.
Repository	Select the repository on the Core where the data from the machines is stored. The repository you select is used for all machines being protected.
Encryption Key	Specify whether encryption is applied to the agent on the machines that is stored in the repository. The encryption key is assigned to all machines that are being protected.
Protection Schedule	Specify 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.
	 NOTE: For more information, see Modifying Protection Schedules .
Initially Pause Protection	Optionally, you can choose to pause protection when first run; that is, the core does not take snapshots of the machines until you manually resume protection.

5. Verify that AppAssure can connect to each machine successfully. To do this, select the check box next to each machine on the **Protect Machines** window, and click **Verify**.
6. The **Protect Machines** window shows an icon next to each machine that reflects its readiness for deployment, as follows:

Icon	Description
Green icon	AppAssure is able to connect to the machine and it is ready to be protected.

Icon	Description
Yellow icon	AppAssure is able to connect to the machine; however, the agent is already paired with a core machine.
Red icon	AppAssure 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.

7. After machines are verified successfully, select the check box next to each machine and click **Protect**.

Monitoring the protection of multiple machines

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

To monitor the protection of multiple machines:

1. Click the **Machines** tab to view the status and progress of the protection.
The **Protected Machines** page appears.
2. Click the **Events** tab to view the related tasks, events, and alerts.
The **Tasks** page appears.

Text Box	Description
To view task information	As volumes are transferred, the status, start times, and end times display in the Tasks pane. Click Details to view more specific information about the task.
To view alert information	As each protected machine is added, an alert is logged that details whether the operation was successful or errors were logged. The level of the alert along with the transactional date and message is displayed. If you want to remove all alerts from the page, click Dismiss All .
To view event information	Details about the machine and the data that is transferred appear in the Events pane. The level of the event, transactional date, and time message appear.

Managing snapshots and recovery points

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

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

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

Viewing recovery points

To view recovery points:

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

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

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

Viewing a specific recovery point

To view a specific recovery point:

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

Info	Description
Actions	The Actions menu includes the following operations that you can perform on the selected recovery point: Mount — Select this option to mount the selected recovery point. For more information about mounting a selected recovery point, see Mounting A Recovery Point For A Windows Machine. Export — From the Export option, you can export the selected recovery point to ESXi, VMware workstation, or HyperV. For more information about exporting selected recovery points, see Exporting Backup Information For Your Windows Machine To A Virtual Machine. Rollback — Select this option to perform a restore from the selected recovery point to a volume you specify. For more information about performing restores from selected recovery points, see Launching A Restore From The AppAssure Core.

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

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

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

Mounting a recovery point for a Windows machine

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

To mount a recovery point for a Windows machine:

1. In the Core Console, do one of the following:
 - Select the **Machines** tab.
 - a. Next to the machine or cluster with the recovery point that you want to mount, select **Mount** from the **Actions** drop-down menu.
 - b. Select a recovery point from the list in the **Mount Recovery Point** dialog box, and then click **Next**. The **Mount Recovery Points** dialog box appears.
 - From the Core Console, select the machine that you want to mount to a local file system.

The **Summary** tab for the selected machine displays.
 - a. Select the **Recovery Points** tab.
 - b. In the list of recovery points, expand the recovery point that you want to mount.
 - c. In the expanded details for that recovery point, click **Mount**.
The **Mount Recovery Points** dialog box appears.
2. In the **Mount** dialog box, edit the text boxes for mounting a recovery point as described in the following table:

Text Box	Description
Mount Location: Local Folder	Specify the path used to access the mounted recovery point.
Volume Images	Specify the volume images that you want to mount.
Mount Type	Specify the way to access data for the mounted recovery point: • Mount Read-only. • Mount Read-only with previous writes. • Mount Writable.

Text Box	Description
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.

Dismounting select recovery points

You can dismount select recovery points that are mounted locally on the Core.

To dismount select recovery points:

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

Dismounting all recovery points

You can dismount all recovery points that are mounted locally on the Core.

To dismount all recovery points:

1. From the Core Console, select 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

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 username and password.
A list is displayed 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: You can also enter the machine ID number in this command instead of the line item number.

A list is displayed that shows the base and incremental recovery points for that machine. This list includes a line item number, date/timestamp, location of volume, size of recovery point, and an ID

number for the volume that includes a sequence number at the end (for example, 293cc667-44b4-48ab-91d8-44bc74252a4f:2), which identifies the recovery point.

8. 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, 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 is: `m 2 23 b /tmp/mount_dir`.



CAUTION: You must not unmount a protected Linux volume manually. 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, you can specify one of the following options:

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



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

To remove recovery points:

1. In the left navigation area of the Core Console, select the machine for which you want to view recovery points, and then click the **Recovery Points** tab.
2. Click the **Actions** menu.
3. Select one of the following options:
 - To delete all currently stored recovery points, click **Delete All**.
 - To delete a set of recovery points in a specific data range, click **Delete Range**. The **Delete** dialog box appears. In the **Delete Range** dialog box, specify the range of recovery points that you want to delete by using a start date and time and an end date and time, and then click **Delete**.

Deleting an orphaned recovery point chain

An orphaned recovery point is an incremental snapshot that is not associated with a base image. Subsequent snapshots continue to build onto this recovery point. Without the base image, the resulting recovery points are incomplete and are unlikely to contain the data needed to complete a recovery. These recovery points are considered to be part of the orphaned recovery point chain. If this situation occurs, the best solution is to delete the chain and crease a new base image.



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

To delete an orphaned recovery point chain:

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



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

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 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 as follows:
 - **Force Snapshot** — Takes an incremental snapshot of data updated since the last snapshot was taken.
 - **Force Base Image** — Takes a complete snapshot of all data on the volumes of the machine.
3. When the notification is displayed in the **Transfer Status** dialog box that the snapshot has been queued, click **OK**.
A progress bar appears next to the machine in the **Machines** tab and displays the progress of the snapshot.

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 Core Console, click the **Machines** tab.
2. Select the machine for which you want to pause protection.
The **Summary** tab for this machine appears.
3. In the **Actions** drop-down menu for that machine, click **Pause**.
4. To resume protection, click **Resume** in the **Actions** menu.

Restoring data

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

If you have replication set up between two cores (source and target), you can only export data from the target core after the initial replication is complete. For details, see [Replicating Agent Data On A Machine](#).

 **NOTE:** 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.

Backup


The backup tab allows you to configure the backup policy and recover your system through the RASR USB key or IDSDM. To use this feature, Windows Backup virtual disk should exist. Windows Backup virtual disk is created during **AppAssure Appliance Configuration Wizard**. For more information see, Rapid Appliance Self Recovery in the *Dell DL43000 Appliance Deployment Guide*. Without a Windows Backup virtual disk, you cannot configure a policy or create Windows backups.

Backup Status

Microsoft Windows backup status is available under the **Last Backup** tab. If a backup is currently running, the information is displayed under the **Current Backup** tab. To view the last backup, perform the following steps:

1. In the Core Console, navigate to the **Appliance** → **Backup** tab.
2. Click the arrow beside the **Status** button to view the status of the backup.
3. The **Last Backup** pane displays the following information:
 - Status
 - State
 - Backup Location
 - Start Time
 - End Time
 - Error Description
 - Items that were backed up

 **NOTE:** The above information is displayed whether the Windows Backup Policy is run or not.



If a backup is running, information regarding **Current Backup Progress** and **Start Time** is displayed.

Windows Backup Policy

To configure a Windows backup policy, perform the following steps:

1. In the Core Console, navigate to **Appliance** → **Backup**.
2. Click the **Configure Policy** button.
The **Windows Backup Policy** window is displayed.
3. Enter the parameters as described below:

Text Box Description

Following items will be backed up:

- OS(C:)
- RECOVERY
- Bare metal recovery
- System State

All of the above are selected by default.

Select the time to schedule the backup:

Enter the time to schedule a backup.

4. Click **Configure**.
Once configured you have the option to **Backup now**, **Delete policy** or **View policy** from the **Windows Backup Policy** window.

About exporting protected data from Windows machines to virtual machines

AppAssure supports both a one-time export or continuous export (to support virtual standby) of Windows backup information to a virtual machine. Exporting your data to a virtual standby machine provides you

with a high availability copy of the data. If a protected machine goes down, you can boot up the virtual machine to then perform recovery.

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




Figure 9. Exporting data to a virtual machine

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

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

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

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

Dynamic and basic volumes support limitations

AppAssure supports taking snapshots of all dynamic and basic volumes. AppAssure also supports 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 cannot be exported. AppAssure has the ability to export complex or non-simple dynamic volumes.

AppAssure version 5.3.1.60393 added a check box in the user interface informing you that exports are restricted to simple dynamic volumes. Before the user interface changed with this version, the option of exporting complex or non-simple dynamic disks would have appeared to have been an option. If you attempted to export these disks, the export job would fail.

Exporting backup information from your Microsoft Windows machine to a virtual machine

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

To export Windows backup information to a virtual machine:

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

The **Select Export Type** dialog box appears.

Exporting Windows data using ESXi export

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

Performing a one-time ESXi export

To perform a one-time ESXi export:

1. In the **Select Export Type** dialog box, click **One-time export**.
2. Click **Next**.
The **ESXi Export - Select Recovery Point** dialog box appears.
3. Select a recovery point to export and then click **Next**.
The **Virtual Standby Recovery Point to VMware vCenter Server/ESXi** dialog box appears.

Defining virtual machine information for performing an ESXi export

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 described as follows:

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

2. Click **Connect**.

Performing a continuous (virtual standby) ESXi export

To perform a continuous (virtual standby) ESXi export:

1. In the **Select Export Type** dialog box, click **Continuous (Virtual Standby)**.
2. Click **Next**.
The **Virtual Standby Recovery Point to VMware vCenter Server/ESXi** dialog box is displayed.
3. Enter the parameters for accessing the virtual machine as described below.

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

4. Click **Connect**.
5. In the **Options** tab, enter the information for the virtual machine as described.

Text Box	Description
Virtual Machine Name	Enter a name for the virtual machine being created. For example, VM-0A1B2C3D4

 **NOTE:** It is recommended to use a name that is derived from the agent name or one that matches the agent name. You can also create a name derived from the hypervisor type, IP address or DNS name.

Memory	Specify the memory usage. You can choose from the following options: • Use the same amount of RAM as source machine • Click Use a specific amount of RAM to specify how much RAM to use. For example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machines. (Recommended)
---------------	--

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.

Version Select the version of the virtual machine.

 **NOTE:** To use vSphere Client to manage virtual machines, select version 8 or earlier.

Resource Pool Enter a name for the resource pool.

6. Click **Start Export**.

Exporting Windows data using VMware workstation export

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

Performing a one-time vmware workstation export

To perform a one-time VMware Workstation export:

1. In the **Select Export Type** dialog box, click **One-time export**.
2. Click **Next**.
The **VM Export - Select Recovery Point** dialog box appears.
3. Select a recovery point to export and then click **Next**.
The **Virtual Standby Recovery Point to VMware Workstation/Server** dialog box appears.

Defining one-time settings for performing a 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 described as follows:

Text Box	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, 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 must enter a valid user name for an account that is registered on the target machine. • If you entered a local path, a user name is not required.
Password	Enter the logon credentials for the virtual machine. • If you specified a network share path, you must enter a valid password for an account that is registered on the target machine. • If you entered a local path, a password is not required.

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 below:

Text Box	Description
Virtual Machine	Enter a name for the virtual machine being created. For example, VM-0A1B2C3D4.

Text Box	Description
	 NOTE: It is recommended to use a name that is derived from the agent name or one that matches the agent name. You can also create a name derived from the hypervisor type, IP address or DNS name.

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. - Click Use a specific amount of RAM to specify how much RAM to use. For example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine. (recommended)
---------------	---

- Click **Export**.

Performing a continuous (virtual standby) VMware workstation export

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

- In the **Select Export Type** dialog box, click **Continuous (Virtual Standby)** and then click **Next**. The **VM Export - Select Recovery Point** dialog box appears.
- Select a recovery point to export and then click **Next**. The **Virtual Standby Recovery Point to VMware Workstation/Server** dialog box appears.
- Enter the parameters for accessing the virtual machine described as follows:

Text Box	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, 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 must enter a valid user name for an account that is registered on the target machine. - If you entered a local path, a user name is not required.
Password	Enter the logon credentials for the virtual machine. - If you specified a network share path, you must enter a valid password for an account that is registered on the target machine. - If you entered a local path, a password is not required.

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

Text Box	Description
Virtual Machine	Enter a name for the virtual machine being created. For example, VM-0A1B2C3D4.  NOTE: It is recommended to use a name that is derived from the agent name or one that matches the agent name. You can also create a name derived from the hypervisor type, IP address or DNS name.
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. - Click Use a specific amount of RAM to specify how much RAM to use; for example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine. (recommended)

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

Exporting Windows data using Hyper-V export

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.

Your DL Appliance supports first-generation Hyper-V export to the following hosts:

- Windows 8
- Windows 8.1
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Your DL Appliance supports second-generation Hyper-V export to the following hosts:

- Windows 8.1
- Windows Server 2012 R2

 **NOTE:** Not all protected machines can be exported to Hyper-V second generation hosts.

Only protected machines with the following Unified Extensible Firmware Interface (UEFI) operating systems support virtual export to Hyper-V second generation hosts:

- Windows 8 (UEFI)
- Windows 8.1 (UEFI)
- Windows Server 2012 (UEFI)
- Windows Server 2012R2 (UEFI)

 **NOTE:** Hyper-V export to second-generation VM can fail if the Hyper-V host does not have enough RAM allocated to perform the export.

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

Performing a one-time Hyper-V export

To perform a one-time Hyper-V export:

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

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

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

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

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

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



NOTE: Hyper-V Export supports VHDX disk formats if the target machine is running Windows 8 (Windows Server 2012) or higher. If the VHDX is not supported for your environment, the option is disabled.

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

Performing a continuous (virtual standby) Hyper-V export

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

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

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

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

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

- **VHD**

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

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

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

Exporting Microsoft Windows data using Oracle VirtualBox export

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

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

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

Performing a one-time Oracle VirtualBox export

Complete the steps in this procedure to perform a one-time export to Oracle VirtualBox.

To perform a one-time Oracle VirtualBox export

1. In the AppAssure Core Console, do one of the following:
 - From the button bar, click **Export** to launch the Export Wizard, and do the following:
 1. On the **Select Export Type** page, select **One-time export**, and then click **Next**.
 2. On the **Protected Machines** page, select the protected machine you want to export to a virtual machine, and then click **Next**.
 - Navigate to the machine you want to export, and then in the **Summary** tab, from the **Actions** drop-down menu for that machine, select **Export > One-time**.

The Export Wizard appears on the **Recovery Points** page.

2. On the **Recovery Points** page, select the recovery point from the AppAssure Core that you want to export, and then click **Next**.
3. On the **Destination** page in the Export Wizard, in the **Recover to Virtual machine** drop-down menu, select **VirtualBox**, and then click **Next**.
4. On the **Virtual Machine Options** page, select **Use Windows machine**.
5. Enter the parameters for accessing the virtual machine as described in the following table.

Option	Description
Virtual Machine Name	Enter a name for the virtual machine being created.
	 NOTE: The default name is the name of the source machine.

Option	Description
Target Path	Specify a local or remote target path to create the virtual machine.  NOTE: The target path should not be a root directory. If you specify a network share path, you will need to enter valid logon credentials (user name and password) for an account that is registered on the target machine. The account must have read and write permissions to the network share.
Memory	Specify the memory usage for the virtual machine by clicking one of the following: - Click Use the same amount of RAM as source machine to specify that the RAM configuration is the same as the source machine. - Click Use a specific amount of RAM to specify how much RAM to use; for example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine. (recommended) 6. To specify a user account for the virtual machine, select Specify the user account for the exported virtual machine, and then enter the following information. This refers to a specific user account for which the virtual machine will be registered in the event there are multiple user accounts on the virtual machine. When this user account is logged on, only this user will see this Virtual Machine in VirtualBox manager. If an account is not specified, then the Virtual Machine will be registered for all existing users on the Windows machine with Oracle VirtualBox. User name - Enter the user name for which the virtual machine is registered. Password - Enter the password for this user account. 7. Click Next. The name that you enter appears in the list of virtual machines in the Hyper-V Manager console. 8. On the Volumes page, select the volume(s) to export. For the virtual machine to be an effective backup of the protected machine include the protected machine's boot drive. Example. C:\. 9. On the Summary page, click Finish to complete the wizard and to start the export.  NOTE: You can monitor the status and progress of the export by viewing the Virtual Standby or Events tab.

Performing a continuous (virtual standby) Oracle VirtualBox export

Complete the steps in this procedure to create a Virtual Standby and perform a continuous export to Oracle VirtualBox.

To perform a continuous (virtual standby) VirtualBox export

- In the AppAssure Core Console, do one of the following:
 - On the **Virtual Standby** tab, click **Add** to launch the Export Wizard. On the **Protected Machines** page of the Export Wizard, select the protected machine you want to export, and then click **Next**.
 - Navigate to the machine you want to export, and, on the **Summary** tab in the **Actions** drop-down menu for that machine, click **Export** > **Virtual Standby**.
- On the **Destination** page in the Export Wizard, in the **Recover to Virtual machine** drop-down menu, select **VirtualBox**, and then click **Next**.
- On the **Virtual Machine Options** page, select **Use Windows machine**.

4. Enter the parameters for accessing the virtual machine as described in the following table.

Option	Description
Virtual Machine Name	Enter a name for the virtual machine being created.  NOTE: It is recommended to use a name that is derived from the agent name or one that matches the agent name. You can also create a name derived from the hypervisor type, IP address or DNS name.
Target Path	Specify a local or remote target path to create the virtual machine.  NOTE: The target path should not be a root directory. If you specify a network share path, you will need to enter valid logon credentials (user name and password) for an account that is registered on the target machine. The account must have read and write permissions to the network share.
Memory	Specify the memory usage for the virtual machine by clicking 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 RAM to use; for example, 4096 MB. The minimum amount allowed is 512 MB and the maximum is determined by the capability and limitations of the host machine. (recommended)

5. To specify a user account for the virtual machine, select **Specify the user account for the exported virtual machine**, and then enter the following information. This refers to a specific user account for which the virtual machine will be registered in the event there are multiple user accounts on the virtual machine. When this user account is logged on, only this user will see this Virtual Machine in VirtualBox manager. If an account is not specified, then the Virtual Machine will be registered for all existing users on the Windows machine with VirtualBox.

- **User name** - Enter the user name for which the virtual machine is registered.
- **Password** - Enter the password for this user account.

6. Select **Perform initial one-time export** to perform the virtual export immediately instead of after the next scheduled snapshot.

7. On the Volumes page, select the volume(s) to export. For the virtual machine to be an effective backup of the protected machine include the protected machine's boot drive. Example. C:.\.

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

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

Virtual Machine Management

The **VM Management** tab displays the status of the protected machines. You can start, stop, and add network adapters (applicable for Hyper-V and ESXi virtual machines only). To navigate to the VM Management tab, click **Appliance** → **VM Management**.

 **NOTE:** The Start, Stop and Add Network Adapter buttons may take up to 30 seconds to appear each time the **Appliance** → **VM Management** tab is selected.

Appliance

Storage

Tasks

Virtual Machine Management

Hyper-V Virtual Standby(s)

Agent / VM Information				Export Status		Hypervisor Information		VM Operations	
Agent Name	VM Name	Status	Location	Status	Last Export	Name	Status		
10.10.101.95	LHyperV-10.10.101.84	Enabled (Running)	C:\VS_SPACE\LocalHV_10.10.101.95\LHyperV-10.10.101.84	Succeeded	3/27/2015 5:02:20 PM	localhost	Online	Start Stop Add Network Adapter	

ESX Virtual Standby(s)

Agent / VM Information				Export Status		Hypervisor Information		VM Operations	
Agent Name	VM Name	Status	Location	Status	Last Export	Name	Status		
10.10.101.84	ESX-10.10.101.84	Disabled (Off)	ESX-10.10.101.84	Failed	4/9/2015 2:45:09 PM	10.10.101.7	Online	Start Stop Add Network Adapter	

Other Virtual Standby(s)

Hypervisor Information		Agent / VM Information		Export Status	
Type	Agent Name	Location	Status	Last Export	
Oracle VirtualBox	Test-10.10.101.96	C:\test	Unknown	Not performed	

VM Management for Hyper-V and ESXi virtual standby(s)

Field	Description
Agent / VM Information	Agent Name: Indicates the name of the protected machine for which you have created virtual standby.
	VM Name: Indicates the name of the VM.
	 NOTE: It is recommended to use a name that is derived from the agent name or one that matches the agent name. You can also create a name derived from the hypervisor type, IP address or DNS name.
Export Status	Status: Indicates the status of the virtual machine. Possible values are- - Running - Stopped - Starting - Suspended - Stopping - Unknown (temporary status)  NOTE: The above status values depend on the hypervisor type. Not all hypervisors display all the status values.
	Location: Indicates the location of VM. For example, D:\export. The VM location must have sufficient space to hold the VM metadata and virtual drives needed for the virtual machine.
	Status - Indicates the following status of an export process: - Complete - Failed - In progress

Field	Description
Hypervisor Information	• Not Performed
	2. If an export is currently in progress, percentage of export is displayed.
	Last Export: Indicates the time of last export.
	Name: Indicates the name of the Hypervisor on which VM is created.
	Status: Indicates the status of connection to the Hyper-V and ESXi hypervisors.
	• Online • Offline • Unknown (temporary status)
	 NOTE: The status is displayed only for Hyper-V and ESXi hypervisors.
VM Operations	Allows you to start or stop the virtual machine, and add a network adapter.

VM Management for Other Virtual Standby(s)

Field	Description
Hypervisor Information	Type: Indicates the type of the Hypervisor.
Agent / VM Information	Agent Name: Indicates the name of the protected machine for which you have created virtual standby. Location: Indicates the location of VM. For example, D:\export. The VM location must have sufficient space to hold the VM metadata and virtual drives needed for the virtual machine.
Export Status	Status 1. Indicates the following status of an export process: • Complete • Failed • In progress • Not performed 2. If an export is currently in progress, percentage of export is displayed as a progress bar. Last Export: Indicates the time of last export.

Creating a virtual network adapter

Virtual machines must have one or more Virtual Network Adapters (VNAs) to connect to the internet. A VM should have a VNA for each real network adapter (RNA) on the protected machine. The VNA and the matching RNA should have a similar configuration. You can add VNAs to your VM when creating the Virtual Standby or you can add VNAs at a later time.

When creating a virtual standby, there is a suggested adapter for every adapter in the protected machine, when configuring a virtual machine. You may add to or remove all or some of these suggested adapters.

The maximum number of VNAs per VM depends on the type of hypervisor. For Hyper-V you can add up to 8 adapters for every virtual machine.

To create a virtual network adapter:

1. Navigate to the **VM Management** page.
2. Click the **Add Network Adapter** button associated with the VM to add a VNA.
 -  **NOTE:** Do not add adapters to a VM for a Virtual Standby that is still running backups or exports of protected machines. The additional VNAs can cause future export operations to fail.
 -  **NOTE:** It is recommend that you add VNAs just before you start the VM in replacement of the protected machine. Ensure that you stop or pause any pending exports for the VM through the virtual standby tab.

The **Virtual Network Adapters and Switches** window appears.

3. Click **Create** to create a virtual network adapter.

The **Create Virtual Network Adapter** window appears.
4. Choose an existing virtual switch from the drop-down menu.
 -  **NOTE:** While selecting virtual switches for ESXi, the dropdown only lists switches with 'VM' or 'Virtual Machine' in their names. Only select a switch of type **Virtual Machine Port Group**, you can verify the type of switch through the ESXi hypervisor GUI.
5. Click **Create**.
 -  **NOTE:** To remove a virtual network adapter, use the hypervisor management interface.

Starting a VM operation

To start a VM operation:

1. Navigate to the **VM Management** window.
2. Click the **Start** button associated with the VM to start.
 -  **NOTE:** The GUI may lag in showing the correct status of the machine. The Start button may remain disabled upto 30 seconds after the buttons have been used. The Start button is enabled only if the virtual machine can be started.
 -  **NOTE:** Do not click the Start button if an export task to the virtual machine is currently running or is likely to start soon. Verify the schedule of the next export task by viewing the **Protected Machines** tab and **Virtual Standby** tab. If an export task has been scheduled in the near future, cancel or skip the export task or, wait for the export task to complete before starting the virtual machine. Exporting data fails if initiated when the virtual machine is running although you can start a virtual machine when an export task is running.
 -  **NOTE:** It is recommended that you do not start the VM that is maintained as a Virtual Standby. Virtual Standby VMs are intended to be active or started as a replacement for a failed protected machine. If the protected machine is still active, first stop or pause any pending exports for the VM through the Virtual Standby tab before starting the VM.

Stopping a VM operation

To stop a VM operation:

1. Navigate to the **VM Management** window.
2. Click the **Stop** button associated with the VM to stop.

 **NOTE:** The Stop button is enabled only if the virtual machine is currently running and is available within a 30 second (approximately) refresh after starting the VM.

 **NOTE:** The Start button is enabled within 30 seconds (approximately) after stopping the VM.

 **NOTE:** Once the protected VM is restored, remove the VM from the hypervisor and its corresponding Virtual Standby. Recreate the Virtual Standby for the restored protected machine. This ensures that the Virtual Standby VM accurately mirrors the protected machine.

Performing a rollback

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

 **NOTE:** Rollback functionality is also supported for your protected Linux machines by using the command-line `aamount` utility. For more information, see [Performing A Rollback For A Linux Machine By Using the Command Line](#).

To perform a rollback:

1. In the Core Console, do one of the following:
 - Click the **Machines** tab, and then do the following:
 - a. In the list of protected machines, select the check box next to the machine that you want to export.
 - b. In the **Actions** drop-down menu for that machine, click **Rollback**.
 - c. In the **Rollback — Select Recovery Point** dialog box, select a recovery point to export, and then click **Next**.
 - In the left navigation area of the AppAssure Core Console, select the machine that you want to roll back, which launches the **Summary** tab for that machine.
 - d. Click the **Recovery Points** tab, and then select a recovery point from the list.
 - e. Expand the details for that recovery point, and then click **Rollback**.
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 recovery point being used for the rollback was created.
Recovery Console Instance	To restore the recovery point to any machine that booted in URC mode, enter the user name and password.

3. Click **Load Volumes**.

The **Volume Mapping** dialog box appears.

 **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** options, select the destination volume to which the selected volume should roll back.

6. Select from the following options:

- **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 point 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 , you can perform a rollback for volumes on your protected Linux machines using the command-line `aamount` utility.



CAUTION: Do not attempt to perform a rollback on the system or root (/) volume.



NOTE: Rollback functionality is supported for your protected Windows machines within the Core Console. For more information, see [Performing A Rollback](#).

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 host name of your AppAssure Core server.

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

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

5. To list the currently mounted recovery points for the specified machine, enter the following command:

```
lr <machine_line_item_number>
```



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

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

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

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

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

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

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

 **NOTE:** It is possible to roll back to `/`, but only when performing a Bare Metal Restore while booted with a Live CD. For more information, see [Performing A Bare Metal Restore For A Linux Machine](#).

7. When prompted to proceed, enter `y` for Yes.
after the rollback proceeds, a series of messages appear that notify you of the status.
8. Upon a successful rollback, the `aamount` utility automatically mounts and reattach the kernel module to the rolled back volume if the target was previously protected and mounted. If not, mount the rollback volume to the local disk and then verify that the files are restored.

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

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

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

About bare metal restore for Windows machines

Servers, when operating as expected, run and perform the tasks they are configured to do. When a catastrophic event occurs, rendering the server inoperable, immediate steps are needed to restore the server to its previous operating condition. The process typically entails reformatting the machine, reinstalling the operating system, recovering data through backups, and reinstalling software applications.

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

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

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

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 server and the functioning Core

- Hardware to restore (new or old, similar or dissimilar)
- Blank CD and CD burning software
- VNC viewer (optional)
- Windows 7 PE (32-bit) Compatible Drivers Storage and Network adapter drivers for the target machine
- Storage Controller, RAID, AHCI, and chipset drivers for the target operating system

 **NOTE:** The Storage Controller Drivers are only needed if the restore being performed is to dissimilar hardware.

Roadmap for performing a bare metal restore for a Windows machine

To perform a BMR for a Windows machine:

1. Create a boot CD. See [Creating A Bootable Cd Iso Image](#).
2. Burn the image to disk.
3. Boot the target server from the boot CD. See [Loading A Boot CD](#).
4. Connect to the recovery disk.
5. Map the volumes. See [Mapping Volumes](#).
6. Initiate the recovery. See [Launching A Restore From The AppAssure Core](#).
7. Monitor the progress. See [Viewing The Recovery Progress](#).

Creating a bootable CD ISO image

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

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

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

To create a bootable CD ISO image:

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

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

Naming the boot cd file and setting the path

To name the boot CD file and set the path:

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

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

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

Creating connections

To create connections:

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

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

Injecting drivers in a boot cd

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

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

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

To inject drivers in a boot CD:

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

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

Creating the boot cd

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

Viewing the iso image creation progress

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

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

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

Accessing the iso image

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

Loading a boot CD

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

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

To load a boot CD:

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

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

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

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

Injecting drivers to your target server

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

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

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

To inject drivers to your target server:

1. Download the drivers from the manufacturer's website for the server and unpack them.
2. Compress the folder that contains the drivers by using a file compressing utility (for example, Win Zip) and copy it to the target server.

3. In the Universal Recovery Console, click **Driver Injection**.
4. To locate the compressed driver file, navigate through the filing system and select the file.
5. If you clicked **Driver Injection** in step 3, click **Add Driver**. If you clicked **Load driver** in step 3, click **Open**.

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

Launching a restore from the Core

To launch a restore from the Core:

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

 **NOTE:** The Host and Password values are the credentials you recorded in the previous task. For more information, see [Loading A Boot CD](#).

8. Click **Load Volumes** to load the target volumes to the new machine.

Mapping volumes

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

For manual mapping, you must already have the new machine correctly formatted before restoring it. For more information, see [Launching A Restore From The AppAssure Core](#).

To map volumes:

1. To automatically map volumes, do the following:
 - 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 both listed beneath and are selected.
 - c. If the destination disk that is automatically mapped is the correct target volume, select **Destination Disk**.
 - d. Click **Rollback**, and then proceed to step 3.
2. To manually map volumes, do the following:
 - a. In the **RollbackURC** dialog box, select the **Manually Map Volumes** tab.
 - b. In the **Volume Mapping** area, under **Source Volume**, verify that the source volume is selected and that the appropriate volumes are both listed beneath and are selected.

- c. Under **Destination**, from the drop-down menu, select the appropriate destination that is the target volume to perform the bare metal restore of the selected recovery point, and then click **Rollback**.
3. In the **RollbackURC** confirmation dialog box, review the mapping of the source of the recovery point and the destination volume for the rollback. To perform the rollback, click **Begin Rollback**.

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

Viewing the recovery progress

To view the recovery progress:

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

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

Starting the restored target server

To start the restored target server:

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

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

Repairing startup problems

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

To repair startup problems:

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

Performing a bare metal restore for a Linux machine

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

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

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

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

 **NOTE:** Before you begin utilizing the AppAssure commands, you can install the screen utility. The screen utility enables you to scroll the screen to view larger amounts of data, such as a list of recovery points. For information about installing the screen utility, see [Installing The Screen Utility](#)

To perform a bare metal restore for a Linux machine:

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

```
sudo aamount
```

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

```
lm
```

5. When prompted, enter the IP address or host name of your AppAssure Core server.
6. Enter the logon credentials, that is, the username and password, for this server.

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

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

```
lr <machine_line_item_number>
```

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

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

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

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

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

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

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

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

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

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

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

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

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

Installing the screen utility

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

To install the screen utility:

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

Creating bootable partitions on a Linux machine

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

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

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

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

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

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

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

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

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

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

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

```
mount --bind /dev/ /mnt/dev

mount --bind /proc/ /mnt/proc

chroot/mnt/bin/bash

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

Viewing events and alerts

To view events and alerts:

1. Do one of the following:
 - In the Core Console, on the Machines tab, click the hyperlink for the machine for which you want to view events.
 - In the left **Navigation** area of the Core Console, select the machine for which you want to view events.
2. Click the **Events** tab.

A log of all events for current tasks and alerts appears.

Protecting server clusters

About server cluster protection

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

You can easily configure an 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 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 is displayed 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 Agent software on each of the machines or nodes in the cluster. AppAssure supports the application versions and cluster configurations listed in the following table.

Table 4. Supported applications and cluster types

Application	Application Version and Related Cluster Configuration	Windows Failover Cluster
Microsoft Exchange	2007 Single Copy Cluster (SCC)	2003, 2008, 2008 R2
	2007 Cluster Continuous Replication (CCR)	
	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
- Dynamic disks

- 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. 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 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, set up at least one repository that is associated with your AppAssure Core.

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

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:

Text Box	Description
Host	The host name or IP address of the cluster, the cluster application, or one of the cluster nodes that you want to protect.  NOTE: If you use the IP address of one of the nodes, this node needs to have an AppAssure agent installed and started.
Port	The port number on the machine on which the AppAssure 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 username.
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 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](#).

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 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. After adding a cluster, navigate to that cluster, and click the **Machines** tab.
2. Click the **Actions** menu, and then click **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 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 for 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 below.

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 Core in which the data from this machine is be stored.
Encryption	Specify whether encryption is to 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 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 information about setting custom intervals, see Protecting A Cluster.

Process of modifying cluster node settings

After 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](#).
3. To configure notification groups for system events, see [Configuring Notification Groups For System Events](#).
4. To customize retention policy settings, see [Customizing Retention Policy Settings](#).
5. To modify the protection schedule, see [Modifying Protection Schedules](#).
6. To modify transfer settings, see [Modifying Transfer Settings](#).

Roadmap for configuring cluster settings

The roadmap for configuring cluster settings involves performing the following tasks:

- Modifying Cluster Settings
- Configuring Cluster Event Notifications
- Modifying the Cluster Retention Policy
- Modifying Cluster Protection Schedules
- Modifying Cluster Transfer Settings

Modifying cluster settings

After adding 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 that you want to modify.
 - In the left navigation area, select the cluster that you want to modify.
2. Click the **Configuration** tab.
The **Settings** page appears.
3. Click **Edit** to modify the settings on this page for the cluster described as follows:.

Text Box	Description
Display Name	Enter a display name for the cluster. The name for this cluster is displayed in the 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 is to 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.

To configure cluster event notifications:

1. Do one of the following:
 - In the Core Console, click the **Machines** tab, and then select the cluster that you want to modify.
 - In the left navigation area, select the cluster that you want to modify.
2. Click the **Configuration** tab, and then click **Events**.

3. Select one of the options described in the following table.

Text Box	Description
Use Core alert settings	This adopts the settings used by the associated core: a. Click Apply. b. Complete Step 5.
Use Custom alert settings	This lets you configure custom settings. Proceed to Step 4.

4. If you select **Custom alert settings**, click **Add Group** to add a new notification group for sending a list of system events.

The **Add Notification Group** dialog box opens.

5. 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 e-mail 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.

6. Click **OK** to save your changes, and then click **Apply**.
7. To edit an existing notification group, next to a notification group in the list, click **Edit**.
The **Edit Notification Group** dialog box appears 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 rollout process that helps in aging and deleting old backups.

1. Do one of the following:
 - In the **Core Console**, click the **Machines** tab, and then select the cluster that you want to modify.
 - In the left navigation area, select the cluster that you want to modify.
2. Click the **Configuration** tab, and then click **Retention Policy**.
3. Select one of the options in the following table:

Text Box	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.

 **NOTE:** If you selected **Custom alert settings**, follow the instructions for setting a custom retention policy as described in [Customizing Retention Policy Settings](#), beginning with Step 4.

Modifying cluster protection schedules

You can modify the protection schedules only if your cluster has shared volumes.

To modify cluster protection schedules:

1. Do one of the following:
 - In the Core Console, click the **Machines** tab, and then select the cluster that you want to modify.
 - In the left navigation area, select the cluster that you want to modify.
2. Click the **Configuration** tab, and then click **Protection Settings**.
3. Follow the instructions for modifying the protection settings as described in [Modifying Protection Schedules](#), starting with Step 2.

Modifying cluster transfer settings

In AppAssure , you can modify the settings to manage the data transfer processes for a protected cluster.

 **NOTE:** You can modify cluster transfer settings only if your cluster has shared volumes.

There are three types of transfers in AppAssure:

Text Box	Description
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 that you want to modify.
 - In the left navigation area, select the cluster that you want to modify.
2. Click the **Configuration** tab, and then click **Transfer Settings**.

3. Modify the protection settings as described in [Modifying Protection Schedules](#), beginning with Step 2.

Converting a protected cluster node to an agent

In AppAssure, 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 want 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 that you want to convert. Click the **Machines** tab for the cluster.
 - From the left navigation area, select the cluster that contains the machine that 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

Viewing cluster system information

To view cluster system information:

1. Do one of the following:
 - In the Core Console, click the **Machines** tab, and then select the cluster that you want to view.
 - In the left **navigation** area, select the cluster that you want to view.
2. Click the **Tools** tab.

The **System Information** page shows 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

For information about viewing events and alerts for an individual machine or node in a cluster, see [Viewing Events And Alerts](#).

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 that you want to view.
 - In the left **Navigation** area, under **Clusters**, select the cluster that you want to view.
2. Click the **Events** tab.

A log displays 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

To view summary information:

1. Do one of the following:
 - In the Core Console, click the **Machines** tab, and then select the cluster that you want to view.
 - In the left **Navigation** area, under **Clusters**, select the cluster that you want 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](#).

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, 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](#).

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 want to view recovery points.
 - In the left navigation area, under **Clusters**, select the cluster for which you want to view recovery points.
2. Click the **Recovery Points** tab.
3. To view detailed information about a specific recovery point, click > next to a recovery point in the list to expand the view.

For information about the operations you can perform on the recovery points, see [Viewing A Specific Recovery Point](#).
4. Select a recovery point to mount.

For information about how to mount a recovery point, see [Mounting A Recovery Point For A Windows Machine](#), starting with Step 2.
5. To delete recovery points, see [Removing Recovery Points](#).

Managing snapshots for a cluster

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 is 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](#). For information about pausing and resuming snapshots for the individual machines in a cluster, see [Pausing And Resuming Protection](#).

Forcing 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 want to view recovery points.
 - In the left navigation area, under **Clusters**, select the cluster for which you want 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

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 want to view recovery points.
 - In the left navigation area, under **Clusters**, select the cluster for which you want 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 as follows:

Text Box	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

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 want to dismount recovery points
 - In the left navigation area, select the cluster for which you want 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

To perform a rollback for SCC (Exchange, SQL) clusters:

1. Turn off all nodes except one.
2. Perform a rollback using the standard AppAssure procedure for the machine as described in [Performing A Rollback](#) and [Performing A Rollback For A Linux Machine By Using the Command Line](#).
3. After the rollback is finished, mount all databases from 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

To perform a rollback for SCC (Exchange, SQL) clusters:

1. Turn off all nodes except one.
2. Perform a rollback using the standard AppAssure procedure for the machine as described in [Performing A Rollback](#) and [Performing A Rollback For A Linux Machine By Using the Command Line](#).
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](#).

Removing 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 that you want to remove
 - In the left navigation area, select the cluster that you want 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](#). To remove a cluster node from protection.

- Do one of the following:
 - In the Core Console, click the **Machines** tab, and then select the cluster that contains the node that you want to remove. On the **Machines** tab for the cluster, select the node that you want to remove.
 - In the left navigation area, under the related cluster, select the node that you want to remove.
- Click the **Actions** drop-down menu and then click **Remove Machine**.
- Select one of the options described in the following table.

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.

Removing all nodes in a cluster from protection

To remove all nodes in a cluster from protection:

- Do one of the following:
 - In the Core Console, click the **Machines** tab, and select the cluster that contains the nodes that you want to remove, then click the **Machines** tab for the cluster.
 - From the left navigation area, select the cluster that contains the nodes that you want to remove, and then click the **Machines** tab.
- Click the **Actions** drop-down menu at the top of the **Machines** tab and then click **Remove Machines**.
- Select one of the options described in the following table.

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.

Viewing a cluster or node report

You can create and view compliance and errors reports about AppAssure activities for your cluster and individual nodes. The reports include AppAssure activity information about the cluster, node, and shared volumes. For more information about AppAssure reporting, see [About Reports](#).

For more information about the exporting and printing options located in the reports toolbar, see [About The Reports Toolbar](#).

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 want to create a report.
 - In the left **Navigation** area, select the cluster or node for which you want 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 Core or AppAssure Agent software 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**.
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.

The report results appear in the page.
6. To export the report results to one of the available format – 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.

Reporting

About reports

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

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

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

About the reports toolbar

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

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

About compliance reports

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

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

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

About errors reports

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

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

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

.

About the Core Summary Report

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

Repositories summary

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

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

- Compression/Dedupe Ratio

Agents summary

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

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

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

Generating a report for a Core or agent

To generate a report for a core or agent:

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

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

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



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

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

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

About the Central Management Console Core reports

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

Generating a report from the Central Management Console

To generate a report from the Central Management Console:

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

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

Completing a full recovery of the DL4000 Appliance

Creating a RAID 1 partition for the operating system

 **CAUTION:** It is essential that you perform these operations only on the RAID 1 virtual disks that contain the operating system. Do not perform these operations on the RAID 6 virtual disks that contain data.

To create a RAID 1 partition:

1. Ensure that the disks in slots 0 and 1 are known working disks.
2. Boot the DL4000 Backup to Disk appliance.
3. When prompted during the boot process, press <Ctrl><R>. The **PERC BIOS Configuration Utility** screen is displayed.
4. Highlight the controller at the top of the **VD Management** tab, and press <F2>, then select **Create New VD**.
 -  **NOTE:** If the RAID-1 OS VD is already present, then fast-init the RAID-1 OS VD.
5. On the **Virtual Disk Management** page, select RAID 1 for RAID Level.
6. Select both disks in the **Physical Disks** box.
 -  **NOTE:** The size of the virtual disk should not exceed 278.87 GB.
7. Enter a VD Name, such as "OS", that identifies the virtual disk as the one that contains the operating system.
8. Press <Tab> to move the cursor to Initialize and press <Enter>.
 -  **NOTE:** The initialization performed at the stage is fast initialization.
9. Click **OK** to finalize the selection or Press <Ctrl><N> twice. The **Ctrl Mgt** page displays.
10. Navigate to the **Select boot device** field and select the virtual disk that contains the operating system. The capacity of this disk is approximately 278 GB.
11. Select **Apply** and press <Enter>.
12. Exit the **PERC BIOS Configuration** utility and press <Ctrl><Alt> to reboot the system.

Installing the operating system

Use the Unified Server Configurator - Lifecycle Controller Enabled (USC-LCE) utility on your appliance to recover the operating system:

1. Locate the operating system installation media.
2. Ensure you have a drive from which to run the media.

You can use a USB optical drive or a virtual media device. Virtual media is supported through iDRAC. For more information on setting up virtual media through iDRAC, see the User Guide for your system's iDRAC device.

If the installation media is corrupt or not readable, then USC may be unable to detect the presence of a supported optical drive. In this case, you may receive an error message stating that no optical drive is available. If the media is not valid (if it is the incorrect CD or DVD, for example), a message displays requesting that you insert the correct installation media.

3. Start USC by booting the system and pressing the <F10> key within 10 seconds of the Dell logo being displayed.
4. Click **OS Deployment** in the left pane.
5. Click **Deploy OS** in the right pane.
6. Select the relevant operating system and click **Next**.

USC extracts the drivers required by the operating system you selected. The drivers are extracted to an internal USB drive named **OEMDRV**.



NOTE: The process for extracting the drivers may take several minutes.



NOTE: All drivers copied by the OS Deployment wizard are removed after 18 hours. You must complete the operating system installation within 18 hours for the copied drivers to be available. To remove the drivers before the 18-hour period is over, reboot the system and press the <F10> key to re-enter USC. Using the <F10> key to cancel the operating system installation or to re-enter USC upon reboot removes the drivers during the 18-hour period.

7. After the drivers are extracted and USC prompts you, insert the operating system installation media.



NOTE: When installing the Microsoft Windows operating system, the extracted drivers are automatically installed during the operating system installation.

Running the recovery and update utility

To run the Recovery and Update Utility:

1. Download the **Recovery and Update Utility** from dell.com/support.
2. Copy the utility to the desktop of the DL4000 Backup to Disk appliance and extract the files.
3. Double-click **launchRUU**.
4. When prompted, click **Yes** to acknowledge that you are not running any of the listed processes.
5. Click **Start** when the **Recovery and update utility** screen displays.
6. When prompted to reboot, click **OK**.

The Windows Server Roles and Features, ASP .NET MVC3, LSI Provider, DL Applications, OpenManage Server Administrator and AppAssure Core Software are installed as part of the Recovery and Update Utility.

7. Reboot your system if prompted again.

8. After all services and applications are installed, click **Proceed**.
The **AppAssure Appliance Recovery** wizard starts.
9. Complete the steps in the **Collecting Information and Configuring** phase of the AppAssure Appliance Recovery Wizard and then click **Next**.
The **Disk Recovery** phase begins.
10. Click **Next** after viewing the warning about AppAssure services being shut down.
The virtual disks for repositories and any virtual standby machines are restored and AppAssure services are restarted. The recovery is complete.

Changing the host name manually

It is recommended that you select a host name during the initial configuration of the DL4000 Backup to Disk Appliance. If you change the hostname at a later time using **Windows System Properties**, you must perform the following steps manually to ensure that the new host name takes effect and the appliance functions properly:

1. Stop AppAssure Core service
2. Delete AppAssure server certificates
3. Delete Core server and registry keys
4. Change the display name in AppAssure
5. Update trusted sites in Internet Explorer

Stopping the Core service

To stop AppAssure Core services:

1. Open **Windows Server Manager**.
2. In the tree on the left, select **Configuration** → **Services**.
3. Right-click **AppAssure Core Service**, and select **Stop**.

Deleting server certificates

To delete AppAssure server certificates:

1. Open a command line interface.
2. Type **Certmgr** and press <Enter>.
3. In the **Certificate Manager** window, select **Trusted Root Certification Authorities** → **Certificates**.
4. Delete any certificate for which the **Issue To** column displays the old host name, and the **Intended Purpose** column displays **Server Authentication**.

Deleting Core server and registry keys

To delete core server and registry keys:

1. Open a command line interface.
2. Type **regedit** and press <Enter> to open the Registry editor.
3. In the tree, navigate to **HKEY_LOCAL_MACHINE** → **SOFTWARE** → **AppRecovery** and open the Core directory.
4. Delete the **webServer** and **serviceHost** directories.

Launching the Core with the new host name

To launch the Core using the new host name that you created manually:

1. Start AppAssure Core services.
2. Right-click the **AppAssure 5 Core** icon on the desktop, and then click **Properties**.
3. Replace the old server name with the new `<server name:8006>`.
For example, **https://<servername>:8006/apprecovery/admin/Core**.
4. Click **OK**, and then launch the AppAssure Core Console by using the **AppAssure 5 Core** icon.

Changing the display name

To change the display name:

1. Log on to the **AppAssure Console** as administrator.
2. Select the **Configuration** tab, and then click the change button on the **General** bar.
3. Enter the new **Display Name** and click **OK**.

Updating trusted sites in Internet Explorer

To update the trusted sites in Internet Explorer:

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

Appendix A— scripting

About powershell scripting

Windows PowerShell is a Microsoft .NET Framework-connected environment designed for administrative automation. AppAssure includes comprehensive client software development kits (SDKs) for PowerShell scripting that enables administrators to automate the administration and management of AppAssure resources by the execution of commands through scripts.

It 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 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, 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 scripts.

Powershell scripting prerequisites

Before using and executing the PowerShell scripts for AppAssure, you must have Windows PowerShell 2.0 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"/>
  </startup>
  <supportedRuntime version="v2.0.50727"/>
</configuration>
```

Testing scripts

If you want to test the scripts you plan to run, you can do so by using the PowerShell graphical editor, `powershell_ise`. You also need to add the configuration file, `powershell_ise.exe.config` to the same folder the configuration file, `powershell.exe.config`.



NOTE: The configuration file, **powershell_ise.exe.config** must have the same content as that of the **powershell.exe.config** file.



CAUTION: If the pre-PowerShell or post-PowerShell script fails, the job also fails.

Input parameters

All available input parameters are used in sample scripts. The parameters are described in the following tables.



NOTE: Script files must possess the same name as the sample script files.

Table 5. AgentTransferConfiguration (namespace Replay.Common.Contracts.Transfer)

Method	Description
<code>public uint MaxConcurrentStreams { get; set; }</code>	Gets or sets the maximum number of concurrent TCP connections the Core establishes to the agent for transferring data.
<code>public uint MaxTransferQueueDepth { get; set; }</code>	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 maximum transfer queue depth.
<code>public uint MaxConcurrentWrites { get; set; }</code>	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.
<code>public ulong MaxSegmentSize { get; set; }</code>	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.
<code>public Priority Priority { get; set; }</code>	Gets or sets the priority for transfer request.
<code>public int MaxRetries { get; set; }</code>	Gets or sets the maximum number of times a failed transfer is retried before it is presumed failed.
<code>public Guid ProviderId { get; set; }</code>	Gets or sets the GUID of the VSS provider to use for snapshots on this host. Administrators typically accept the default.
<code>public Collection<ExcludedWriter>ExcludedWriterIds { get; set; }</code>	Gets or sets the collection of VSS writer IDs, which is 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.

Method	Description
<code>public ushort TransferDataServerPort { get; set; }</code>	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 uses the port number specified in the <code>BlockHashesUri</code> and <code>BlockDataUri</code> properties of the <code>VolumeSnapshotInfo</code> object for each snapped volume.
<code>public TimeSpan SnapshotTimeout { get; set; }</code>	Gets or sets the amount of time to wait for a VSS snapshot operation to complete before giving up and timing out.
<code>public TimeSpan TransferTimeout { get; set; }</code>	Gets or sets the amount of time to wait for further contact from the Core before abandoning the snapshot.
<code>public TimeSpan NetworkReadTimeout { get; set; }</code>	Gets or sets the timeout for network read operations related to this transfer.
<code>public TimeSpan NetworkWriteTimeout { get; set; }</code>	Gets or sets the timeout for network write operations related to this transfer.

Table 6. BackgroundJobRequest (namespace `Replay.Core.Contracts.BackgroundJobs`)

Method	Description
<code>public Guid AgentId { get; set; }</code>	Gets or sets the ID of the agent.
<code>public bool IsNightlyJob { get; set; }</code>	Gets or sets the value indicating whether the background job is a nightly job.
<code>public virtual bool InvolvesAgentId(Guid agentId)</code>	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`.

Method	Description
<code>public uint RamInMegabytes { get; set; }</code>	Gets or sets the memory size for the exported VM. Set to zero (0) to use the memory size of the source machine.
<code>public VirtualMachineLocation Location { get; set; }</code>	Gets or sets the target location for this export. This is an abstract base class.

Method	Description
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)

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.

Method	Description
public VolumeNameCollection VolumeNames { get; set; }	Gets or sets the collection of volume names for transfer.
public ShadowCopyType ShadowCopyType { get; set; }	Gets or sets the type of copying for transfer. Available values: Unknown, Copy, and 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.
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.

Table 7. TransferPostscriptParameter (namespace Replay.Common.Contracts.PowerShellExecution)

Method	Description
<code>public VolumeNameCollection VolumeNames { get; set; }</code>	Gets or sets the collection of volume names for transfer.
<code>public ShadowCopyType ShadowCopyType { get; set; }</code>	Gets or sets the type of copying for transfer. Available values: Unknown, Copy, and Full.
<code>public AgentTransferConfiguration TransferConfiguration { get; set; }</code>	Gets or sets the transfer configuration.
<code>public AgentProtectionStorageConfiguration StorageConfiguration { get; set; }</code>	Gets or sets the storage configuration.
<code>public string Key { get; set; }</code>	Generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.
<code>public bool ForceBaseImage { get; set; }</code>	Gets or sets the value indicating whether the base image was forced.
<code>public bool IsLogTruncation { get; set; }</code>	Gets or sets the value indicating whether the job is a log truncation.
<code>public uint LatestEpochSeenByCore { get; set; }</code>	Gets or sets latest epoch value.
<code>public Guid SnapshotSetId { get; set; }</code>	Gets or sets the GUID assigned by VSS to this snapshot.
<code>public VolumeSnapshotInfoDictionary VolumeSnapshots { get; set; }</code>	Gets or sets the collection of snapshot info for each volume included in the snap.

Table 8. TransferPrescriptParameter (namespace Replay.Common.Contracts.PowerShellExecution)

Method	Description
<code>public VolumeNameCollection VolumeNames { get; set; }</code>	Gets or sets the collection of volume names for transfer.
<code>public ShadowCopyType ShadowCopyType { get; set; }</code>	Gets or sets the type of copying for transfer. Available values: Unknown, Copy, and Full.
<code>public AgentTransferConfiguration TransferConfiguration { get; set; }</code>	Gets or sets the transfer configuration.
<code>public AgentProtectionStorageConfiguration StorageConfiguration { get; set; }</code>	Gets or sets the storage configuration.
<code>public string Key { get; set; }</code>	Generates a pseudorandom (but not cryptographically secure) key, which can be used as a one-time password to authenticate transfer requests.

Method	Description
<code>public bool ForceBaseImage { get; set; }</code>	Gets or sets the value indicating whether the base image was forced.
<code>public bool IsLogTruncation { get; set; }</code>	Gets or sets the value indicating whether the job is a log truncation.
<code>public uint LatestEpochSeenByCore { get; set; }</code>	Gets or sets latest epoch value.

Table 9. VirtualMachineLocation (namespace `Replay.Common.Contracts.Virtualization`)

Method	Description
<code>public string Description { get; set; }</code>	Gets or sets a human-readable description of this location.
<code>public string Method { get; set; }</code>	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>`.

Table 10. VolumeName (namespace `Replay.Common.Contracts.Metadata.Storage`)

Method	Description
<code>public string GuidName { get; set; }</code>	Gets or sets the ID of the volume.
<code>public string DisplayName { get; set; }</code>	Gets or sets the name of the volume.
<code>public string UrlEncode()</code>	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.

<code>public string GetMountName()</code>	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>`.

Method	Description
<code>public override bool Equals(object obj)</code>	Determines whether this instance and a specified object, which must also be a <code>VolumeNameCollection</code> object, have the same value. (Overrides <code>Object.Equals(Object)</code> .)
<code>public override int GetHashCode()</code>	Returns the hash code for this <code>VolumeNameCollection</code> . (Overrides <code>Object.GetHashCode()</code> .)

Table 11. VolumeSnapshotInfo (namespace `Replay.Common.Contracts.Transfer`)

Method	Description
<code>public Uri BlockHashesUri { get; set; }</code>	Gets or sets the URI at which the MD5 hashes of volume blocks can be read.
<code>public Uri BlockDataUri { get; set; }</code>	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>`.

Pretransferscript.ps1

The **PreTransferScript** is executed on the agent side prior to transferring a snapshot.

```
# 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\
\AppRecovery 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.TransferConfiguration

    echo 'StorageConfiguration:'
    $TransferPrescriptParameterObject.StorageConfiguration
}
```

Posttransferscript.ps1

The **PostTransferScript** is executed on the agent side after transferring a snapshot.

```
# 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\
\AppRecovery 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.

```
# 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\
\AppRecovery 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 contains this script in the PowerShell script folder as **PostExportScript.ps1**.

```
# 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\
\AppRecovery 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 nightly job on Core side. It has **\$JobClassName** parameter, that helps to handle those child jobs separately.

```
# 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\
\AppRecovery 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 nightly job on Core side. It has **\$JobClassName** parameter, that helps to handle those child jobs separately.

```

# 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
\AppRecovery 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;
}
}

```

Sample 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

Getting help

Finding documentation and software updates

In the AppAssure Core console there are direct links to AppAssure, Appliance documentation, and software updates. To access the links, click the **Appliance** tab, and then click **Overall Status**. Links to the software updates and documentation are located under the **Documentation** section.

Contacting Dell

 **NOTE:** If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer-service issues, go to software.dell.com/support.

Documentation feedback

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