
An SQL Server DBA's Guide to Actifio Copy Data Management

Copyright, Trademarks, and other Legal Matter

Copyright © 2021 Google LLC. All rights reserved.

Actifio™, OnVault™, and VDP™ are trademarks of Google LLC.

All other brands, product names, goods and/or services mentioned herein are trademarks or property of their respective owners.

Contents

Preface	vii
The ActifioNOW Customer Portal.....	vii
Actifio Support Centers.....	vii
Chapter 1 - Introduction	1
Actifio Data Virtualization.....	1
Capturing SQL Server Data.....	2
Capture Mechanisms	3
The Actifio Connector.....	3
Capturing Microsoft SQL Server Data.....	3
Capturing SQL Server Database Logs.....	4
Resizing a Database Log's Staging VDisk.....	4
SQL Server Data Capture Options	5
Replicating SQL Server Data.....	6
Replicating Logs.....	6
Accessing SQL Server Data.....	7
Workflows to Automate Access to SQL Server Data	8
Actifio Virtual Data Pipeline Working with Existing Backup Products.....	9
Chapter 2 - Required SQL Server Roles for the Windows User	11
Windows Local Admin User	11
Required SQL Roles for the Windows User.....	11
Credentials for Capturing SQL Server Database Logs.....	12
Credentials for Restoring a Microsoft SQL Server Database	12
Credentials for Mounting an SQL Server Database as a Virtual Application	12
Chapter 3 - Adding a SQL Server Database Host & Discovering the Database	13
Before You Begin	13
Adding the Host.....	13
Discovering SQL Server Instances	15
Finding the Discovered Instances and Databases in the App Manager.....	16
Chapter 4 - Capturing Microsoft SQL Server Instances and Databases	17

Before You Begin.....	17
Capturing a Microsoft SQL Server Database.....	20
Configuring Application Settings for Microsoft SQL Server Databases.....	22
Database Log Protection in an SLA Policy.....	23
Configuring Advanced Settings: Policy Settings Overrides.....	24
Chapter 5 - Mounting a Microsoft SQL Server Database	29
Mounting Captured Microsoft SQL Data.....	29
Mounting an SQL Server Database as a New Virtual Database.....	31
Mounting Encrypted SQL Data.....	34
Determining if SQL TDE is Enabled.....	34
Troubleshooting SQL Server Encryption.....	36
SQL Server Master Key, Encryption Certificate, and Password Procedures.....	37
Chapter 6 - Mounting Databases into SQL AlwaysOn Availability Groups	39
Creating an SQL Server AAG in an Actifio Snapshot Pool.....	39
Creating an SQL Server AAG Outside of An Actifio Snapshot Pool.....	39
Creating the New SQL Server AlwaysOn Availability Group.....	40
Chapter 7 - Mounting and Migrating SQL Data	41
Step 1: Mount or Restore.....	42
Step 2: Scheduling the Migration.....	44
Step 3: Finalize	47
Chapter 8 - Cloning SQL Server Databases	49
Chapter 9 - Restoring SQL Server Databases	53
Microsoft SQL Server Database Restore Overview.....	54
Restoring Microsoft SQL Instances and Databases.....	55
Restoring a SQL Server Database to a Different Host.....	56
Restoring SQL Server Databases in a Consistency Group.....	56
Restoring SQL System Databases.....	57
Restoring to an SQL Server Cluster	58
Chapter 10 - Restoring Members of an SQL AlwaysOn Availability Group	59
Identifying the Last Known Good Image of the SQL Server Database.....	59
Restoring the Database on the Primary AAG Node.....	59
Synchronizing Secondary Databases to the Restored Primary Database.....	60
Recovering the Primary From a Non-Corrupt Local Secondary.....	60
Restoring a Secondary SQL Server Database From an Actifio Mirror Copy.....	60
Restoring a Secondary SQL Server Database From an Actifio Dedup DR Copy	61
Rebuilding the SQL AlwaysOn Availability Group.....	62
Error Messages.....	63

Preface

The information presented in this guide is intended for users who are familiar with basic Actifio processes and procedures as described in **Getting Started with Actifio Copy Data Management** and are qualified to administer Microsoft SQL Server databases.

Note: Only users qualified to administer Microsoft SQL Server databases should attempt the procedures presented here. Procedures attempted by unqualified personnel can result in data loss.

Your Actifio appliance's Documentation Library contains detailed, step-by-step, on how to protect and access your data. Each guide is in PDF format and may be viewed online, downloaded, or printed on demand.

The ActifioNOW Customer Portal

During the configuration and initialization of your Actifio appliance your Actifio representative provided you with a user name and password for the ActifioNOW customer portal.

From the customer portal you can obtain detailed reports about your Actifio appliance as well as search the portal's knowledge base for answers to specific questions.

To log into the ActifioNOW customer portal:

1. Go to: <https://now.actifio.com>
2. When prompted, enter the user name and password provided by your Actifio representative.

Actifio Support Centers

To contact an Actifio support representative, you can:

- Send email to: support@actifio.com
- Call:

From anywhere: +1.315.261.7501

US Toll-Free: +1.855.392.6810

Australia: 0011 800-16165656

Germany: 00 800-16165656

New Zealand: 00 800-16165656

UK: 0 800-0155019

1 Introduction

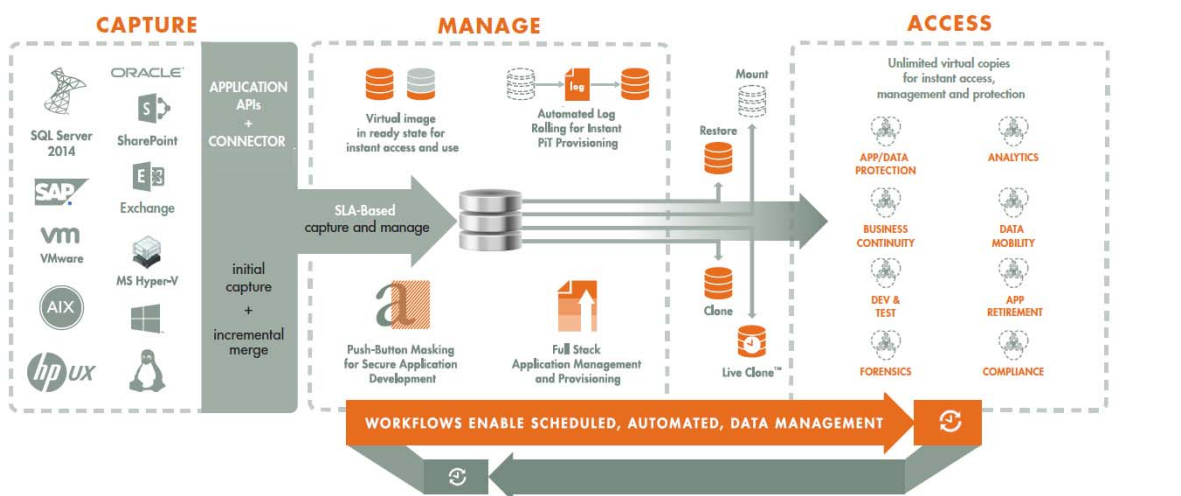
This chapter provides an overview of Actifio concepts and procedures:

- [Actifio Data Virtualization](#) on page 1
- [Capturing SQL Server Data](#) on page 2
- [Capture Mechanisms](#) on page 3
- [Capturing Microsoft SQL Server Data](#) on page 3
- [Capturing SQL Server Database Logs](#) on page 4
- [Replicating Logs](#) on page 6
- [Resizing a Database Log's Staging VDisk](#) on page 4
- [SQL Server Data Capture Options](#) on page 5
- [Replicating SQL Server Data](#) on page 6
- [Accessing SQL Server Data](#) on page 7
- [Workflows to Automate Access to SQL Server Data](#) on page 8
- [Actifio Virtual Data Pipeline Working with Existing Backup Products](#) on page 9

Actifio Data Virtualization

An Actifio Appliance is a highly scalable copy data management platform that virtualizes application data to improve the resiliency, agility, and cloud mobility of your business. It works by virtualizing data in much the same way other technologies have virtualized servers and networks.

Actifio VDP enables users to capture data from production systems, manage it in the most efficient way possible, and use virtual or physical copies of the data whenever and wherever they are needed.



Application data is captured at the block level, in application native format, according to a specified service level agreement (SLA). A “Golden copy” of that data is created and stored once, and is then updated incrementally with only the changed blocks of data in an “incremental forever” model. Unlimited virtual copies of the data can then be made available instantly for use, without proliferating physical copies and taking up additional storage infrastructure.

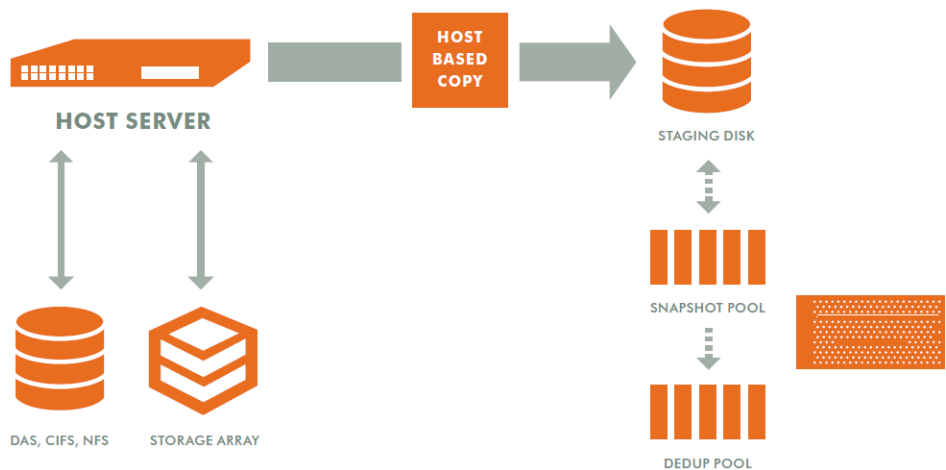
Capturing SQL Server Data

The Actifio user interfaces allow you to capture SQL:

- Instances
- Databases in AlwaysOn Availability Groups
- Consistency Groups of databases
- Individual databases
- System Databases
- User databases
- Databases in VMs.

Actifio VDP moves and manages the Microsoft SQL Server data separately from where Microsoft SQL Server writes its primary storage.

An Actifio Appliance stores application data on a staging disk. Snapshots on the staging disk allow the Actifio Appliance to maintain historical data.



Actifio Data Capture

When capturing data:

1. A staging disk is automatically created and mounted on a server.
2. An initial full copy is made to the staging disk. Subsequent copies consist only of changed blocks.
3. The staging disk is unmounted from the server.
4. A snapshot of the staging disk is made on the Actifio Appliance.

Actifio offers an alternative configuration where production data storage is controlled by an Actifio Appliance. With this approach snapshots and changed-block tracking are native to the production storage array. This approach to data management is known as External Snapshot Pools.

Capture Mechanisms

Customers with the most highly transactional databases are able to take Actifio snapshots of their databases without a performance impact to their users. Actifio customers routinely run SQL Server snapshots every 3 hours (even during the day) on databases hosting airline reservation systems without any negative user impact. Other customers use Actifio to snapshot their Exchange DAG primary nodes without triggering failovers.

VDP uses the built-in Microsoft VSS provider. This comes with safeguards for both performance and space management that Microsoft has built, and is in-line with Microsoft's best practices. Actifio's approach eliminates the performance impact encountered with VMware and third-party backup tool-initiated VSS snapshots, and subsequently uses a patented CBT to efficiently capture database data, including SQL Server and Exchange databases, from the native VSS snapshot.

The Actifio Connector

Actifio VDP captures data by making an initial full copy of the data, then copying only incremental changes. This capability requires the ability to track the changes that occur between capture operations. To track those changes, VDP uses the Actifio Connector.

The Actifio Connector is used to capture data with granularity at the Microsoft SQL Server database level. The Actifio Connector is a small-footprint, lightweight service that can be installed on either virtual or physical servers. The Actifio Connector makes use of Microsoft SQL Server VSS Writer (SqlServerWriter) for discovery, capture, and access operations. SqlServerWriter is installed by default with Microsoft SQL Server.

The Actifio Connector allows you to capture instances, AAG databases, consistency groups of databases, entire Microsoft SQL Server instances, and selected databases in an instance. It also offers options for handling individual Microsoft SQL Server database transaction logs. In addition, it allows you to capture databases that cannot be snapped by VMware without introducing a performance impact.

Specifically, the Actifio Connector:

- Discovers Microsoft SQL Server databases.
- Captures a database by first quiescing, then capturing, then releasing the database. For consistency groups and instances of databases, members are quiesced and released together there by ensuring a consistent point in time capture of data.
- Identifies changes to database data for Actifio's incremental forever capture strategy.
- Captures and manages transaction logs:
 - o Captures Microsoft SQL Server database(s) and logs with one SLA
 - o Truncates Microsoft SQL Server database transaction logs
 - o Rolls Microsoft SQL Server database transaction logs forward for point-in-time recovery when accessing virtual copies.
- Captures databases on VMware VMs, even on pRDMs, avoiding virtual server "stun" issues.

Capturing Microsoft SQL Server Data

Capturing Microsoft SQL Server data consists of four steps:

1. Add servers that host Microsoft SQL Server databases.
2. Discover VMs and Microsoft SQL Server databases.
3. Define Actifio Policy Templates and Resource Profiles according to your RPOs and RTOs. Databases that use the Microsoft SQL Server Full Recovery Model can capture both the database and its logs, so a captured database can be recovered to a point in time by rolling its logs forward.
4. Assign Actifio Policy Templates and Resource Profiles to Microsoft SQL Server databases.

Capturing SQL Server Database Logs

Database log capture is set in a Snapshot policy's Details & Settings. It enables a single Snapshot policy to capture logs for Microsoft SQL Server databases and consistency groups that contain Microsoft SQL Server databases.

The frequency with which database logs are captured is defined separately from that of the database. For example, a database can be captured every day and its logs captured every hour.

The frequency of database log backup is set in minutes, and the frequency at which logs are captured must not exceed the frequency at which its associated database is captured. For example, if a database capture frequency is every 24 hours, the log file capture frequency must be equal to or less than every 24 hours.

Log retention is also defined separately from its associated database. Having separate retention rates allows you to maintain enough log information to cover all Snapshot, Dedup, and OnVault versions of a database. For example, if a database's Snapshot data is kept for three days and its Dedup data kept for seven days, you can define log retention to span all seven days. In this example, a single captured database image can be selected and its logs can be rolled forwards over the seven day period.

Database logs are not deduplicated, and regardless of how many logs are captured during a specified log retention period, a database's captured logs are staged to a single VDisk in the Actifio Snapshot pool. To conserve space in the Snapshot pool, you can use an advanced setting to instruct the database to compress its logs.

You can specify to replicate Microsoft SQL Server database transaction logs to a remote Actifio Appliance. You can use the logs at the remote site for any database image within the retention range of the replicated logs.

Resizing a Database Log's Staging VDisk

The physical space required to accommodate a database's logs is automatically managed by VDP. At a minimum, VDP evaluates typical log sizes and their retention period and add space as needed.

To more efficiently and effectively manage the storage requirements for a database's logs, Snapshot policies provide the following advanced settings:

- **Log Backup Retention Period:** Log retention is defined separately from its associated database. Having separate retention rates allows you to maintain enough log information to cover all Snapshot and Dedup versions of a database. The log retention period is a mandatory setting.
- **Log Staging Disk Size Growth:** Defines the percent at which to automatically grow the staging VDisk on which the logs reside.
- **Estimated Change Rate:** Defines the daily change (%), which allows the Actifio Appliance to better calculate the size of the staging disk needed to hold logs.
- **Compress Database Log Backup:** Instructs the source database to compress its logs before capture on the Actifio Appliance. The database server performs log compression during log backup (default is Enabled).

SQL Server Data Capture Options

When capturing Microsoft SQL Server data, you have the capability of:

- [Capturing Instances, Individual Databases, and Groups of Databases](#) on page 5
- [Capturing Consistency Groups](#) on page 5
- [Capturing a VM's Databases and Boot Volume](#) on page 5

Capturing Instances, Individual Databases, and Groups of Databases

The Actifio Connector is used to capture instances, user databases, system databases, and groups of databases on physical and virtual servers.

When capturing an SQL Instance, you have the option of capturing the entire instance or selected databases within the instance. When you protect the entire instance, as databases are added to the instance, they will automatically be included in the next Actifio capture job. Databases in an instance are quiesced and captured together with a single Actifio SLA.

If Actifio database and log capture are enabled on the SLA Policy, then all databases in that instance can be recovered to the same point-in-time. Recovery and rolling forward of the logs for all or individual databases in an instance is performed from the Actifio user interface with a single action.

Individual members of an instance can be accessed by mount, clone, LiveClone, and restore operations as needed.

Capturing Consistency Groups

A consistency group is a group of databases that are quiesced and captured together with a single Actifio SLA Policy Template and Resource Profile. Membership to a consistency group is assigned manually and is suitable to groups of databases whose members do not change very often. To automatically protect new members of a group of databases, create and protect those databases in an SQL Instance instead.

As the name implies, consistency groups ensure consistent point-in-time capture and recovery across multiple databases. If Actifio's database and log capture technology is enabled on the SLA Policy, then all databases in that group can be recovered to the same point-in-time. Recovery and rolling forward of the logs for all or individual databases in a consistency group is performed from the Actifio user interface with a single action. Members of a consistency group must reside in the same instance.

A consistency group can be made up of:

- One or more system databases
- One or more user databases
- System and/or user databases together
- Zero or more file systems (drive letters or mount points)

Individual members of a consistency group can be accessed by mount, clone, LiveClone, and restore operations.

Databases in a clustered failover instance must be discovered from the active node. Once protected, VDP follows the active SQL node in a cluster. Protection jobs continue to run even in a failover condition. In addition to making capture and access operations easy and fast, consistency groups consume fewer system resources (VDisks) than protecting databases individually.

You can validate the integrity of database backup periodically by mounting a backup image to a server and running database consistency check. You can use the Workflow feature to automate the validation process.

Capturing a VM's Databases and Boot Volume

When capturing databases on VMs you have the option of also capturing the VM's boot volume. When a VM's boot volume is captured along with its databases, an image can be presented that is a fully functional database and VM. The image can then be migrated to a new, permanent location.

Replicating SQL Server Data

Data can be replicated to a second Actifio Appliance or to the cloud for recovery, disaster recovery, or test/dev purposes. Data replication has traditionally been an inhibitor to efficient data management in a geographically distributed environment. Actifio replication addresses these issues with global deduplication and compression that:

- Drives down overall network usage.
- Eliminates the need for a dedicated WAN accelerator/optimizer.
- Does not require storage array vendor licenses as data is sent from one Actifio Appliance to another.
- Is heterogeneous between supported arrays: Tier 1 to Tier 2 and/or Vendor A to Vendor B.
- Encrypts data using the AES-256 encryption standard. Authentication between Actifio Appliances is performed using 1024-bit certificates.

Replication is controlled by Actifio Policy Template policies:

- Production to Mirror policies have several options to replicate data to a second Actifio Appliance.
- Dedup Backup to Dedup DR policies use an Actifio-proprietary replication engine to replicate data to a second Actifio Appliance. In addition Dedup Backup to Dedup DR policies allow you to replicate data to two locations.
- Production to OnVault policies use an Actifio proprietary engine to transfer data to object storage.

Replicating Logs

When a policy's **Enable Database Log Backup** is set to **Enable**, the Replicate Logs advanced setting allows Microsoft SQL Server database transaction logs to be replicated to a remote Actifio Appliance. For a log replication job to run, there must be a replication policy (StreamSnap, Dedup-Async, or Remote Dedup) included in the template along with a resource profile that specifies a remote Actifio Appliance, and at least one successful replication of the database must first be completed. You can then use the logs at the remote site for any database image within the retention range of the replicated logs. This function is enabled by default.

Log replication uses StreamSnap technology to perform the replication between the local and remote Actifio Appliances; log replication goes directly from the local snapshot pool to the snapshot pool on the remote appliance.

Accessing SQL Server Data

For Microsoft SQL Server databases that use the Full Recovery Model, VDP can instantly present a copy of the database rolled forward to a specific point of time. The roll forward operation is performed from AGM.

For Microsoft SQL Server databases that use the Simple Recovery Model, VDP can instantly present the most recent backup of the database.

Regardless of the Microsoft SQL Server recovery model used, Microsoft SQL Server data can be accessed via a Fibre Channel or iSCSI interface, just as if accessing a traditional storage system.

Role-based Access Control

You can control which users have access to data, Actifio features, and resources. Captured data can be marked sensitive, and Actifio users can be granted access permission to sensitive data.

Mounts

The Actifio mount function provides instant access to data without moving data. Captured copies of databases can be rolled forward via the Actifio user interface and mounted on any database server.

Actifio VDP provides two ways to mount a Microsoft SQL Server database:

- **The Virtual Application mount** presents and makes the captured Microsoft SQL Server data available to a target server as an Microsoft SQL Server database. This allows you to create and manage copies of production databases for non-production use. Virtual application mounts are created from the Actifio Appliance and do not require manual intervention by database, server, or storage administrators. Virtual application mounts can be used for database reporting, analytics, integrity testing, and test and development. Virtual databases are detailed in [Mounting an SQL Server Database as a New Virtual Database](#) on page 31 and [Chapter 6, Mounting Databases into SQL AlwaysOn Availability Groups](#).
- **The standard mount**, also called a direct mount, presents and makes the captured Microsoft SQL Server data available to a target server as a file system, not as a database. This is useful if a database is corrupt, lost, or if a database server is being replaced. In such cases you cannot use a restore operation to recover the database. Instead, you can mount an image and copy the database files from the mounted image to their original location on the database server. Direct Mounts are detailed in [Mounting Captured Microsoft SQL Data](#) on page 29

LiveClones

A LiveClone is an independent copy of Microsoft SQL Server data that can be refreshed and masked before being made available to users. This enables development and test teams to work on the latest set of data without having to manually manage the data or interfere with the production environment.

Clones

The clone function moves a copy of the production data to a different location from the source. The amount of time required to complete a clone operation depends on the amount of data involved. Clones are detailed in [Chapter 8, Cloning SQL Server Databases](#).

Restores

A restore reverts the production data to a specified point in time. Restore operations actually move data. Restore operations are typically performed after a massive data corruption or storage array failure. The amount of time required to complete a restore operation depends on the amount of data involved.

To restore a database and then apply logs, the restored database must be in Restoring Mode. You can restore the database in Restoring Mode and then roll the logs forward to a specific point in time. If you restore the database without specifying Restore with no Recovery, the database will be restored and brought online without applying logs. Restores are detailed in [Chapter 9, Restoring SQL Server Databases](#) and in [Chapter 10, Restoring Members of an SQL AlwaysOn Availability Group](#). For a near-zero-downtime restore, mount the data first as detailed in [Chapter 7, Mounting and Migrating SQL Data](#).

Workflows to Automate Access to SQL Server Data

Workflows automate access to the captured Microsoft SQL Server data. Workflows can present data as a direct mount or as a LiveClone:

- Direct mounts (standard or application aware) work well for Microsoft SQL Server data that does not need to be masked prior to being presented. A mounted copy of data can be refreshed manually or automatically on a schedule. Direct mounts allow you to instantly access captured Microsoft SQL Server data without actually moving the data.
- A LiveClone is a copy of your production Microsoft SQL Server data that can be updated manually or on a scheduled basis. You can mask sensitive data in a LiveClone prior to making it available to users.

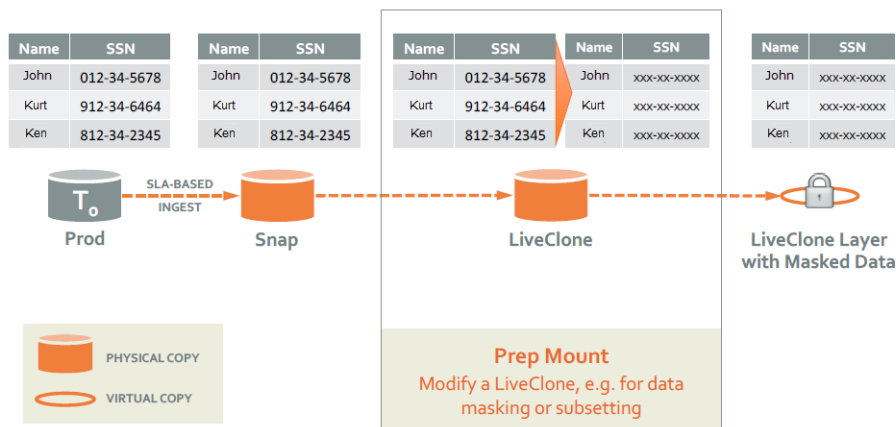
Combining Actifio's automated Microsoft SQL Server data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Users can provision their own environments almost instantly.

For example, an Actifio administrator can create an SLA Template Policy that captures Microsoft SQL Server data according to a specified schedule. The administrator can mark the captured production Microsoft SQL Server data as sensitive and only accessible by users with the proper access rights.

After access rights have been defined and data has been captured, the administrator can create a Workflow that:

- Makes the captured Microsoft SQL Server data available as a LiveClone or as a direct mount.
- Updates the LiveClone or mountable Microsoft SQL Server data on a scheduled or on-demand basis
- Optionally automatically applies scripts to the LiveClone's Microsoft SQL Server data after each update. This is useful for masking sensitive Microsoft SQL Server data.

Once the Workflow completes, users with proper access can provision their environments with the LiveClone or mountable Microsoft SQL Server data.



Workflow With Masked Social Security Data

For more information, refer to [Creating Automated Workflows for SQL Server Databases](#).

Actifio Virtual Data Pipeline Working with Existing Backup Products

As more and more enterprises look to speed up the application development using production databases, Actifio VDP is often required to coexist with legacy backup products working off the same production database environments. Actifio can perfectly co-exist with other products capturing data from production databases, if these best practices are followed.

Actifio has a proprietary method of Change Block Tracking (CBT) so backup solutions using native SQL or other methods of obtaining the backups are not impacted by a scheduled Actifio data capture jobs.

Traditional backup jobs can be very I/O intensive. They may have long durations, and may impact performance of the database during the backup windows. Actifio has made strides to minimize impact during jobs, but even a block-level incremental-forever update must generate some I/O, and must take a little time.

Requirement	Do not schedule legacy backup software and Actifio to run jobs in a way that allows any overlap in time.
Best Practice	Schedule Actifio database jobs to begin at a time when the legacy backup software should be finished. Do not schedule the legacy backup software to run immediately after an Actifio job would normally complete.
Reason	If legacy backup jobs and Actifio jobs run concurrently, it may result a serious performance impact on the database server leading to instability and possibly an outage.

Database logs are used to capture individual transactions in a database, enabling point-in-time recoveries. Most agility use cases center around getting database snapshots on a periodic basis from production. Common frequency ranges from daily to weekly or once in two weeks, depending on the use case. As a result, application developers do not commonly have the need to position their non-prod instance to a specific point-in-time from the source (production). This usually eliminates the need to capture and manage logs as a part of an Actifio agility solution.

Requirement	Only one system can manage (capture and/or truncate/purge) logs, either the legacy backup software or Actifio.
Best Practice	Continue to allow all log management be performed by the legacy backup software, do not use Actifio to protect logs in this environment.
Reason	If Actifio VDP is configured to manage (capture and/or truncate/purge) logs, and the legacy backup software is also capturing and/or truncating/purging logs, then one or both systems may end up with an incomplete log chain, making it difficult or impossible to recover the database to a specific point in time.

2 Required SQL Server Roles for the Windows User

Microsoft SQL Server requires specific user roles to perform specific operations. To perform Actifio capture, restore, unmount, delete, and Virtual Application mount operations on an SQL Server database, you must provide VDP with credentials for a Windows user (a local user or a domain user) who has been assigned a role with sufficient SQL privileges to perform the operation.

This chapter details the user roles required to perform capture, restore, unmount, delete, and Virtual Application mount operations from an Actifio Appliance. The recommended roles presented in this chapter are based on Microsoft's best practices for accessing SQL Server databases.

Note: *Creating users and assigning roles must be done by qualified system and database administrators. If users are improperly defined, and/or roles are improperly applied, the result can lead to Actifio job failure, security violations, and possible data loss.*

This chapter describes permissions associated with:

- [Windows Local Admin User](#) on page 11
- [Required SQL Roles for the Windows User](#) on page 11
- [Credentials for Capturing SQL Server Database Logs](#) on page 12
- [Credentials for Restoring a Microsoft SQL Server Database](#) on page 12
- [Credentials for Mounting an SQL Server Database as a Virtual Application](#) on page 12

Windows Local Admin User

To perform capture, restore, unmount delete, and Virtual Application Mounts, the Actifio Connector must be installed with the credentials of a Microsoft Windows user who has sufficient privileges in the SQL environment. The Windows user must be assigned a specific role or roles. The Microsoft Windows user can be a newly created or existing user.

Required SQL Roles for the Windows User

A Windows Local Admin user assigned to the `sysadmin` server role will have all necessary permissions to perform Actifio capture, restore and Virtual Application Mounts.

If the `sysadmin` server role is deemed too liberal, then assign a Windows user the following roles:

- `dbcreator` server role
- `db_backupoperator` database role
- `db_owner` database role

In addition, such users must also be assigned the following securables:

- `View any database`

- Create any database
- Alter any database
- Connect SQL

The following sections detail where to enter the Windows Local Admin's username and password to perform specific Actifio SQL related operations.

Note: In the following procedures, when entering user names, in most cases the domain name and user name (domain\username) format will be sufficient. In rare cases, entering the domainname\username will return the error: Logon failure: unknown user name or bad password [1326] In such cases, use the fully qualified domain name format: (username@fqdn) to address the problem.

Credentials for Capturing SQL Server Database Logs

When applying an Actifio SLA Policy Template to an SQL Server database, if the template contains a policy that captures database logs you must enter credentials of a Windows user assigned the proper role(s) in the AGM in the application's SLA Application Details & Settings.

Note: Credentials are required for logs; they are not required if only databases are being captured.

Credentials for Restoring a Microsoft SQL Server Database

When restoring SQL Server databases from AGM, in the Restore dialog box, enter credentials of a Windows user assigned the proper role(s).

Credentials for Mounting an SQL Server Database as a Virtual Application

A Virtual Application Mount mounts an SQL Server database as a virtual application. When performing a Virtual Application Mount of an SQL Server database from an Actifio Appliance, the user must be assigned a role that allows both the ability to mount and unmount (detach) the SQL Server database.

When performing a Virtual Application Mount, in the Mount dialog box Advanced Options, enter the credentials of a Windows user assigned the proper role(s).

3 Adding a SQL Server Database Host & Discovering the Database

Before You Begin

Before you can protect SQL Server databases:

- Review your network configuration, including firewall ports, as detailed in **Network Administrator's Guide to Actifio VDP**. Pay special attention to *Notes on Discovering Specific Microsoft Application Types*.
- Install the Actifio Connector on the database hosts, also detailed in **Network Administrator's Guide to Actifio VDP**.
- If your databases are in VMware VMs, be sure to review **A VMware vCenter Administrator's Guide to Actifio Copy Data Management**.
- Make sure the database permissions are set correctly; see [Chapter 2, Required SQL Server Roles for the Windows User](#).

Adding a SQL Server Database Host and Discovering the Database

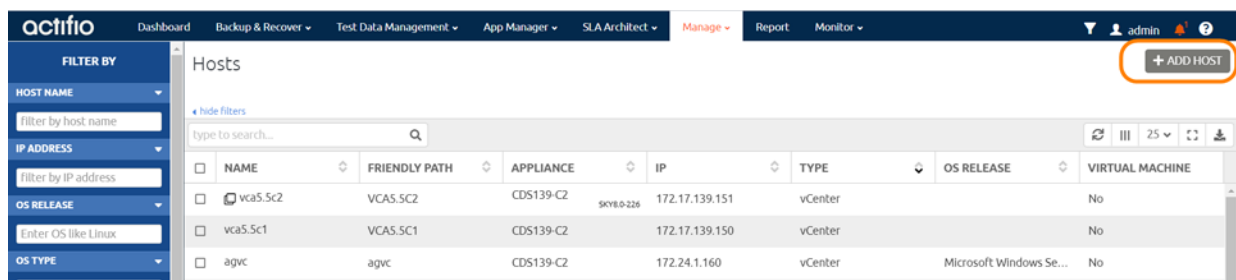
Before you can protect a SQL Server database, you must add the host and discover the database. This requires:

1. [Adding the Host](#) on page 13
2. [Discovering SQL Server Instances](#) on page 15
3. [Finding the Discovered Instances and Databases in the App Manager](#) on page 16

Adding the Host

Add the host to AGM. If the host is already added then edit the host and make sure to set all the configurations correctly.

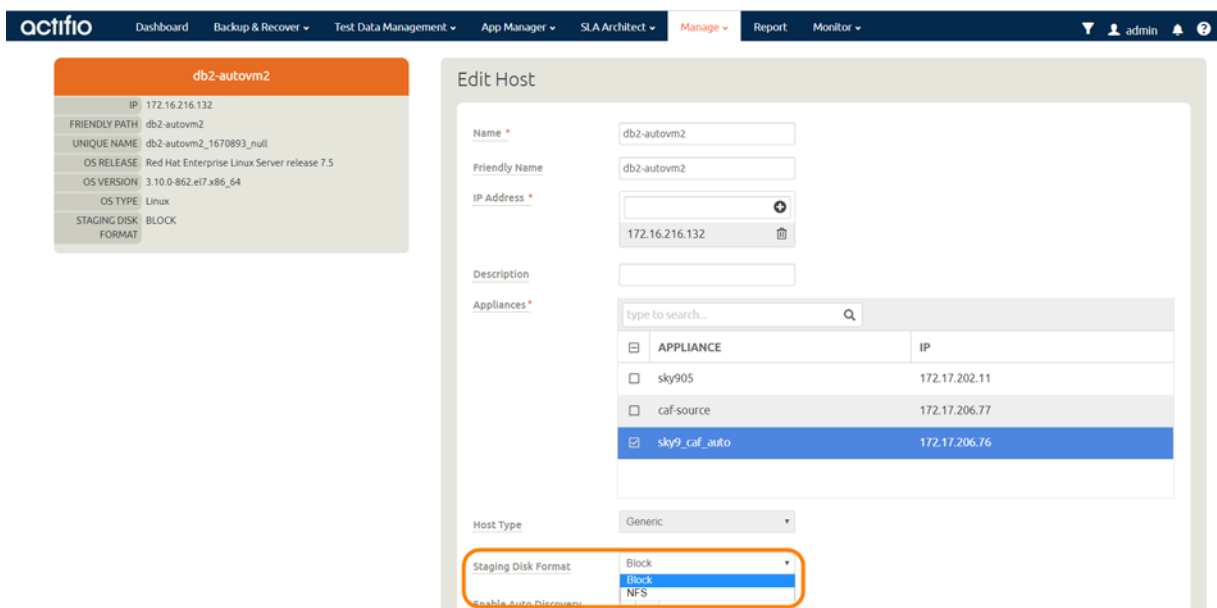
1. From the Manage, Hosts list, click **+Add Host**.



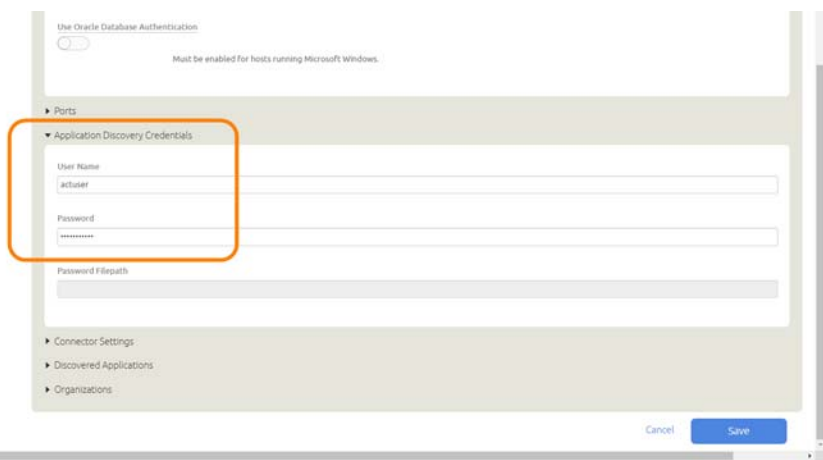
The screenshot shows the Actifio App Manager interface. The top navigation bar includes 'Dashboard', 'Backup & Recover', 'Test Data Management', 'App Manager', 'SLA Architect', 'Manage', 'Report', and 'Monitor'. The 'Manage' dropdown is active, showing 'Hosts'. On the right side of the 'Manage' dropdown, there is a '+ ADD HOST' button highlighted with a red circle. The main content area displays a table of hosts with columns: NAME, FRIENDLY PATH, APPLIANCE, IP, TYPE, OS RELEASE, and VIRTUAL MACHINE. The table contains three rows of data.

NAME	FRIENDLY PATH	APPLIANCE	IP	TYPE	OS RELEASE	VIRTUAL MACHINE
vca5.5c2	VCA5.5C2	CDS139-C2	172.17.139.151	vCenter		No
vca5.5c1	VCA5.5C1	CDS139-C2	172.17.139.150	vCenter		No
agvc	agvc	CDS139-C2	172.24.1.160	vCenter	Microsoft Windows Se...	No

2. On the Add Host page:
 - o **Name:** Provide the database server name.
 - o **IP Address:** Provide the database server IP and click the + sign on the right corner.
 - o **Appliances:** Select the check box for the appliance.
 - o **Host Type:** Make sure this is Generic.
3. Click **Add** at bottom right to add the host.
The host is added.
4. Right-click the host and select **Edit**.
5. On the Edit Host page: Select the **Staging Disk Format:**
 - o For block-based backup with CBT or GPFS: select **Block**
 - o For file-based backup with Full+Incremental file system backup: select **Block or NFS**



6. In Application Discovery Credentials, enter the username/password that you set up in [Before You Begin](#) on page 13.

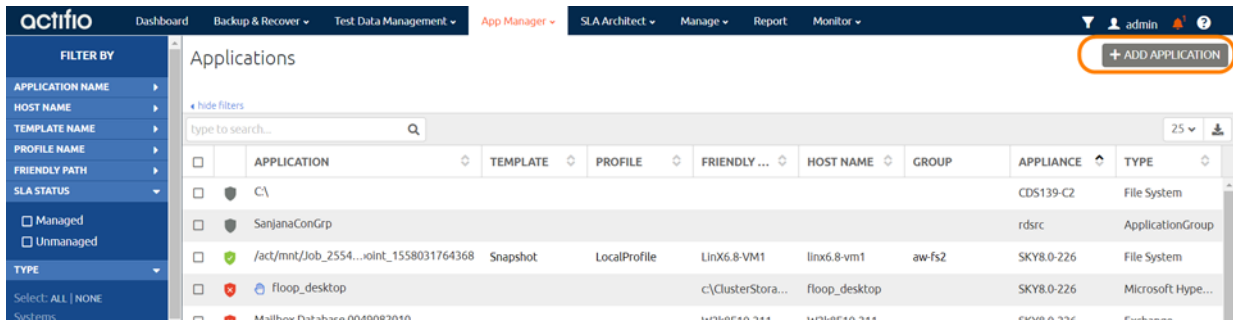


7. Select **Save** at the bottom of Edit Host page.

Discovering SQL Server Instances


To discover SQL Server instances:

1. From the App Manager, Applications list, select **Add Application** in the upper right corner.



2. On the Add Application page, select **SQL Server**, then select the SQL Server database host. If you have many hosts, you can use the search feature or use the filter to see only hosts that are managed by a specific Actifio Appliance.

Note: When adding SQL instances that participate in Microsoft Server Failover Clusters, either:
* Run the wizard once for each node in the cluster,
or else
* Add "All applications" and multi-select all nodes at once. With this method, you must then apply protection from the Applications listing view: select "Manage SLA" for the desired SQL Instance.



3. Select the host and click **Next** in the bottom right corner. This will run the discovery on the SQL Server host and will discover all SQL Server instances and databases running on it.

actifio Dashboard Backup & Recover Test Data Management App Manager SLA Architect Manage Report Monitor admin

Microsoft SQL Server Onboarding

1 Discover 2 Select 3 Manage 4 Configure 5 Finish

Select the host on which you wish to discover Microsoft SQL Server Instances / AGs

[+ Add Host](#)

HOST	IP ADDRESS	FRIENDLY PATH	APPLIANCE
W12S12STD.winqa.actifio.com	172.27.74.11	172.27.74.11	SKYRESSQLONB
W12S14STD1.sqlqa.actifio.com	172.27.75.13	172.27.75.13	SKYRESSQLONB
W12S16STD.sqa.actifio.com	172.27.74.13	172.27.74.13	SKYRESSQLONB
W16S14STD.winqa.actifio.com	172.27.74.15	172.27.74.15	SKYRESSQLONB
W16S19STD.winqa.actifio.com	172.27.74.14	172.27.74.14	SKYRESSQLONB

[Cancel](#) [Next](#)

Finding the Discovered Instances and Databases in the App Manager

To find the newly-discovered database, go to the App Manager Applications list. All applications known to the AGM of all types are listed. Use the Type application filter on the left pane to show only SQL Server database instances. The new SQL Server instances and databases will appear in the list as unmanaged (the red shield icon).

actifio Dashboard Backup & Recover Test Data Management App Manager SLA Architect Manage Report Monitor admin

Applications

[+ ADD APPLICATION](#)

hide filters SLA: Unmanaged Application Type: SQL Server AG Application Type: SQL Server Database Application Type: SQL Server Instance

type to search...

	APPLICATION	TEMPLATE	PROFILE	FRIENDLY PATH	HOST NAME	APPLIANCE	TYPE
<input type="checkbox"/>	AAG_GRP10			W2K14MSCS2456...	W2K14MSCS2456...	SKYRESSQLONB	SQL Server AG
<input type="checkbox"/>	AAG_GRP20			W2K14MSCS2456...	W2K14MSCS2456...	SKYRESSQLONB	SQL Server AG
<input type="checkbox"/>	AG12CLU11			MS12AGCLU11.WI...	MS12AGCLU11.WI...	SKYRESSQLONB	SQL Server AG
<input type="checkbox"/>	AG12CLU11-Grp1			MS12AGCLU11.WI...	MS12AGCLU11.WI...	SKYRESSQLONB	SQL Server AG
<input type="checkbox"/>	AGS14_Grp45			MS12AGCLU3.SQA...	MS12AGCLU3.SQA...	SKYRESSQLONB	SQL Server AG
<input type="checkbox"/>	AG_DB1			MS16AGCLU23.SQ...	MS16AGCLU23.SQ...	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	AG_DB2			MS16AGCLU23.SQ...	MS16AGCLU23.SQ...	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	AG_DB3			MS16AGCLU23.SQ...	MS16AGCLU23.SQ...	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	ActifioDB1			W16S14AG23	W16S14AG23	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	ActifioDB10			W16S14AG23	W16S14AG23	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	ActifioDB11			W2K14MSCS2456...	W2K14MSCS2456...	SKYRESSQLONB	SQL Server Database
<input type="checkbox"/>	ActifioDB12			W2K14MSCS2456...	W2K14MSCS2456...	SKYRESSQLONB	SQL Server Database

1 - 500 of 1758 applications 1 of 4 pages

4 Capturing Microsoft SQL Server Instances and Databases

This chapter includes:

[Before You Begin](#) on page 17

[Capturing a Microsoft SQL Server Database](#) on page 20

[Configuring Application Settings for Microsoft SQL Server Databases](#) on page 22

[Database Log Protection in an SLA Policy](#) on page 23

[Configuring Advanced Settings: Policy Settings Overrides](#) on page 24

AGM allows you to capture:

- Instances
- Primary database of an AlwaysOn Availability Group
- Consistency Groups
- Individual members of a Consistency Group
- System Databases
- User databases
- Databases in VMs

Before You Begin

Before you can protect SQL Server databases:

- Add the hosts and discover their databases using the AGM SQL Server wizard under Backup & Recover or Test Data Management as detailed in [Chapter 3, Adding a SQL Server Database Host & Discovering the Database](#).
- Create SLA Policy Templates and Resource Profiles that define how to protect the databases. Procedures for creating SLA Policy templates and resource profiles are in the AGM online help.

Capturing Databases in an Instance vs a Consistency Group

A database is quiesced, then captured, then released. For consistency groups and instances of databases, members are quiesced and released together for a consistent point-in-time capture of data.

When capturing an **SQL Instance**, as databases are added to the instance, they will automatically be included in the Actifio capture operation. Capturing databases in an SQL instance lends itself to environments where databases are regularly added and removed. Databases mounted to an SQL Instance as virtual applications are not protected with the other members of the instance. Virtually mounted databases must be protected separately.

Membership to a **consistency group** is done manually. Capturing databases in a Consistency Group lends itself to environments where databases are not often added or removed.

Out-of-Band Data

When capturing Microsoft SQL Server data, production data is typically controlled by a non-Actifio storage controller on your existing storage arrays. AGM moves and manages the Microsoft SQL Server data separately from where Microsoft SQL Server writes its primary storage.

Databases that use the Microsoft SQL Server Full Recovery Model can take advantage of Actifio's database and log capture technology. This technology allows you to define a single policy to capture both the database and its logs. Because a single policy captures both the database and its logs, a captured database can be recovered to a point in time by rolling its logs forward via the Actifio appliance's user interface.

When capturing data:

- A staging disk is automatically created and mounted on a server via Fibre Channel or iSCSI.
- An initial full copy is made to the staging disk. Subsequent copies consist only of incremental change blocks.
- The staging disk is unmounted from the server.
- A snapshot of the staging disk is made on the Actifio appliance.

Protecting SQL Server Databases in IBM Storwize or Pure Storage FlashArray Storage

If out-of-band Microsoft SQL Server databases are in volumes on external storage pools, put all system databases (model, master, and msdb) on a different volume used for storing databases. This prevents them from being overwritten when databases on the volume are restored. See *Adding an External Storage Array* in the AGM online help.

In-Band (CDS only)

Actifio offers an alternate configuration where production data storage is controlled by a managed Actifio appliance. With this approach snapshots and changed-block tracking are native to the Actifio appliance and the Actifio appliance is placed in the data path between the SAN and the application host.

If Microsoft SQL Server databases are on in-band volumes, Actifio recommends that you put all system databases (model, master, msdb, and tempdb) on a different volume used for storing databases. This prevents them from being overwritten when databases on the in-band volume are restored.

Note: *In-band management is provided only by Actifio CDS appliances. Actifio Sky and CDX appliances do not support in-band operations.*

In-Band Data, Out-of-Band Data, and VM Management

Microsoft SQL servers are managed differently whether they are in-band or out-of-band, or as part of an entire VM.

Managed as an Application, In-Band or Out-of-Band ESP	Managed as an Application, Out-of-Band	Managed as Part of a VM
<p>Actifio appliances manage the entire volume(s) that the database resides on. Restore operation restores the entire volume.</p> <p>Note: Use caution when restoring an in-band SQL server database as it can overwrite data used by other applications.</p>	<p>Actifio appliances manages only the database files.</p>	<p>Entire VMware VMs are managed using VMware APIs; entire Hyper-V VMs are managed using the Actifio Connector.</p> <p>Note: If you are managing SQL databases that are part of an entire managed VM, see A VMware Administrator's Guide to Actifio Copy Data Management.</p>
<p>Actifio Connector coordinates the VSS snapshot and performs log truncation.</p>		<p>The VMware API coordinates the VSS Snapshot. The Actifio Connector must be installed on the VM for log truncation.</p>
<p>Placing the LUN in-band uses Actifio change block tracking.</p>	<p>The Actifio Connector uses change block tracking on named files (very efficient for large database files).</p>	<p>The VMware API provides change block tracking. For Hyper-V, the Actifio Connector provides change block tracking.</p>
<p>Transaction logs are backed up when a backup job runs if you select Truncate Log After Backup in Details & Settings (see Configuring Advanced Settings: Policy Settings Overrides on page 24).</p>		<p>Transaction logs are not captured.</p>
<p>Client can roll forward with logs.</p>		<p>Roll forward not supported during restore.</p>

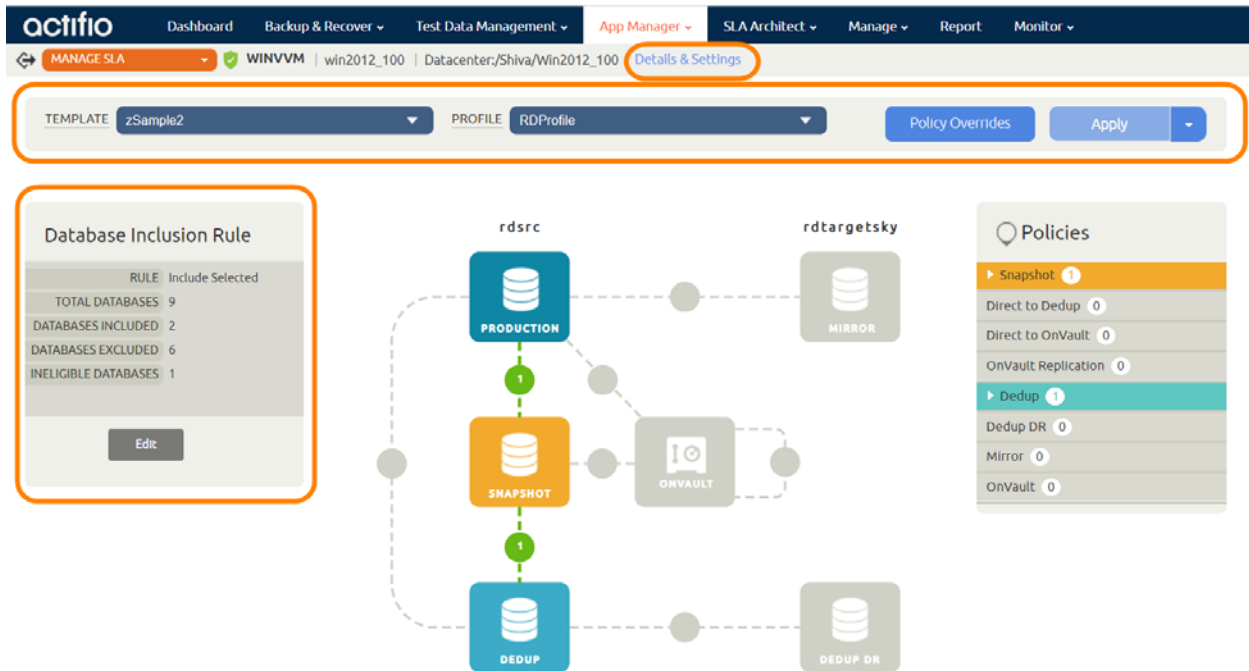
Note: Microsoft SharePoint data can be managed by capturing its Microsoft SQL Server database. When capturing a Microsoft SharePoint SQL Server database, application settings specific to SharePoint are listed.

Note: If a failover cluster and a standalone cluster reside in the same location, the databases in these clusters could be protected twice if databases in the AAG and on a **failover** instance are discovered as part of the **instance**, or if databases in the AAG and on a **stand-alone** instance are discovered as part of the **AAG**.

Capturing a Microsoft SQL Server Database

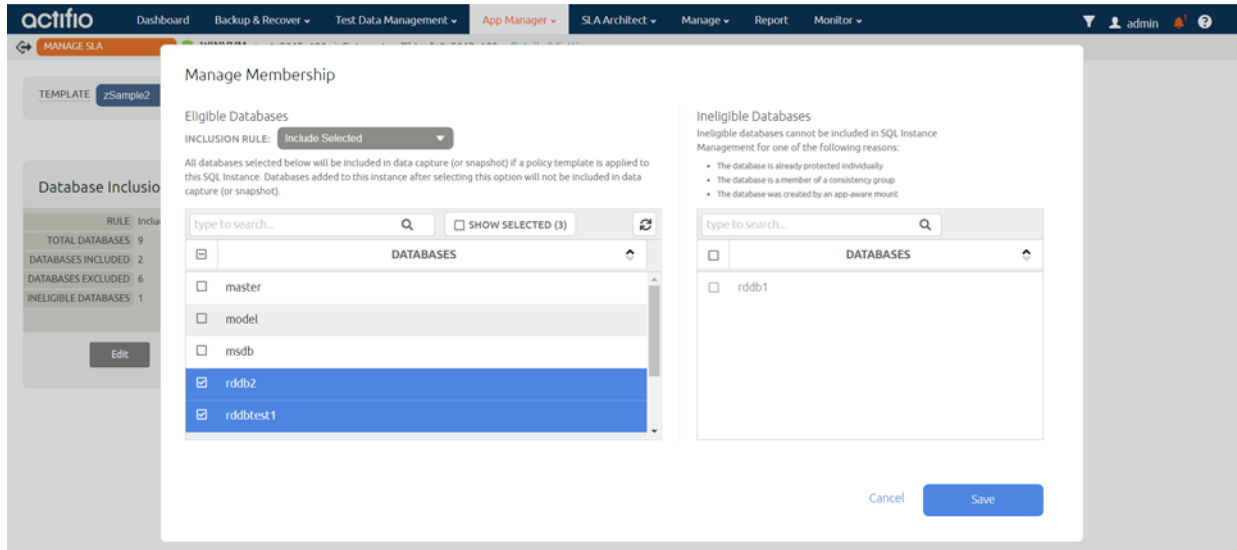
To capture a Microsoft SQL Server database:

1. From the AGM navigation, go to App Manager > **Applications**. The Applications page opens.
2. Right-click the Microsoft SQL Server database, instance, AAG, or consistency group that you want to capture and select **Manage SLA**. The Manage SLA page opens.



The Manage SLA Page

3. From the Manage SLA window, choose a Template and Profile from the drop-down lists:
 - o **Template:** An existing SLA template that includes policies to define the snapshot/deduplication/replication of the application data.
 - o **Profile:** An existing SLA resource profile that defines the resources used to store the data of the application as snapshot/deduplicated/replicated images.
4. From the Manage SLA Template window make the following changes prior to applying an SLA:
 - o **Details and Settings:** Settings specific to Microsoft SQL such as application type, host name, host IP address, path, operating system, VDP appliance, and appliance IP address. See [Configuring Application Settings for Microsoft SQL Server Databases](#) on page 22 for details.
 - o **Policy Overrides:** Override specific policy settings previously configured in the selected SLA template. Policy Overrides can be useful or required in certain circumstances. You can only override policy settings if the policy's template has been configured to allow policy settings overrides, See [Configuring Advanced Settings: Policy Settings Overrides](#) on page 24 for details.
5. To select databases, under Database Inclusion Rule, click **Edit**. The Manage Membership dialog box opens:



Managing Membership

6. From the Manage Membership dialog box, select the databases to capture.
7. Click **Save** and the Manage Membership dialog box closes.
8. Click **Apply** to apply the SLA template and resource profile and the success message box appears.

The selected database(s) are not captured until the scheduled job runs according to the hours of operations defined in the SLA template. For example, if at 10:00 am you assign a template that has hours of operation from 2:00 am to 5:00 am, then the first job will not start until the VDP appliance has an available job slot after 2:00 am.

Configuring Application Settings for Microsoft SQL Server Databases

From the Application Details & Settings dialog box (accessed through the application's **Details & Settings**), you can modify application-specific settings for capturing Microsoft SQL Server databases. Application settings may be useful or required in certain circumstances. After you configure your application settings, click **Save Changes**.

Note: To reset one or more application settings back to its default state, click **Reset to Default** to the right of the selection you want to reset.

Note: This table details all of the SQL application details and settings. The actual list displayed depends on whether an SQL Instance, SQL Cluster, or AAG is selected.

Table 1: SQL Server Application Details & Settings

Application Setting	Description
Username/Password	User credentials needed for backing up database transaction logs. This account must have backup privileges. Credentials are required only if you select Truncate Log or Backup Transaction Log and the local system does not have permissions to the SQL Server database. See Chapter 2, Required SQL Server Roles for the Windows User for details on roles and permissions.
Staging Disk Size (GB)	Enter the staging disk size in the Staging Disk Size (GB) field: 1 to 256000. The Connector calculates the maximum size of the database as configured and adds 20%. The Staging Disk Size option allows you to allocate a staging disk to hold backup and to allow future growth of the database.
Staging Disk Mount Point	Use this to define where to mount the staging disk.
SQL Database Backup Path	Enter an SQL Database Backup Path to define a location for a temporary SQL backup. If the VDP Connector needs to take a full, native backup of the SQL Server database, that backup will be saved in this directory. Ensure there is enough free space on the volume hosting this directory to hold a full database backup. Note: This setting is only applicable for a VM that includes the Actifio Connector and is configured for a Microsoft SQL Server database.
Service Access Point IP Address (SQL server availability groups only)	Enter a value here to back up from a SQL availability appliance. Specify the IP address of the appliance node you want the database to be backed up from. This option is not required if you want the database to be backed up from the active node and it is not required for a failover appliance.
Use Service IP for Restore	Honor the service access point IP during restore for a clustered application.
Connector Options	Leave Connector Options blank unless directed to enter a value by Actifio Support.
Log Staging Disk Size	Enter a Log Staging Disk Size (GB) to override the space automatically defined for database log backups. Valid entries are 1 to 4000.

Database Log Protection in an SLA Policy

When creating a snapshot policy for a database you can also capture its log files. The frequency at which database logs are captured is defined separately from that of the database. For example, a database can be captured every day and its logs captured every hour. The frequency of database log backup is set in minutes, and the frequency at which logs are captured must not exceed the frequency at which its associated database is captured. For example, if a database is captured every 24 hours, the log file capture frequency must be less than every 24 hours.

Frequency and retention are defined in the Details & Settings of the database snapshot policy. Log capture is done without regard to when its associated database is captured.

You enable the Log Protection functionality through the Enable Database Log Backup advanced settings in an SLA snapshot policy. Frequency and retention are defined in the Details & Settings for an SLA Policy.

Policy Settings

No
 Yes

ENABLE DATABASE LOG BACKUP

RPO (MINUTES) * 15
Must have a minimum value of 15 and a maximum of 1440.

LOG BACKUP RETENTION PERIOD (IN DAYS) * 5
Must have a minimum value of 1 and a maximum of 90.

REPLICATE LOGS (USES STREAMSNAP TECHNOLOGY)
 Yes
 No

LOG STAGING DISK GROWTH SIZE (IN PERCENT)
Must have a minimum value of 5 and a maximum of 100.

ESTIMATED CHANGE RATE
Must have a minimum value of 0 and a maximum of 100.

COMPRESS DATABASE LOG BACKUP
 Yes
 No

JOB BEHAVIOR WHEN TARGET VM NEEDS SNAPSHOT CONSOLIDATION
 Fail the job if VM needs consolidation

Cancel Save Changes

Policy Settings

The physical space required to accommodate a database's logs is automatically managed by AGM. At a minimum, AGM will evaluate typical log sizes and their retention period and add space as needed. To more efficiently manage the storage requirements for a database's logs, Snapshot policies provide the following advanced settings:

- **Log Backup Retention Period:** Log retention is defined separately from the retention of the Snapshot policy. Having a separate retention period allows you to use logs in conjunction with copies of the database stored in both the Snapshot and Dedup pools. The log retention period is a mandatory setting.
- **Log Staging Disk Size Growth:** Defines the percent at which to automatically grow the staging VDisk on which the logs reside. This setting is from 5 to 100 percent.
- **Estimated Change Rate:** Defines the daily change (in percent), which allows the VDP appliance to better calculate the size of the staging disk needed to hold logs. This setting is from 0 to 100.
- **Compress Database Log Backup:** Instructs the source database to compress its logs before capture by VDP. The database server performs log compression during log backup.

You can replicate database logs to a remote Actifio Appliance, and use the remote logs for any database image within the retention range of the replicated logs. Log replication uses StreamSnap technology between the local and remote appliances; log replication goes directly from the local snapshot pool to the snapshot pool on the remote appliance. For a log replication job to run, there must be a replication policy (StreamSnap, Dedup-Async, or Remote Dedup) included in the template, and at least one successful replication of the database must first be completed.

Configuring Advanced Settings: Policy Settings Overrides

Click **Policy Overrides** in the Manage SLA window to show the Policy Settings Override dialog. From here you can override specific policy settings associated with the selected SLA template. After you are done, click **Save Changes**.

Note: You can override policy settings in the Application Manager only if the policy template **Allow Overrides on Policy Settings** parameter has been set to Yes.

Note: To reset a policy override setting to its default state, click the check box to the left of the selection; click **Select options that will revert back to default** to reset all policy override settings back to their default state.

Table 2: Policy Settings Overrides Valid for SQL Server Instances, AAG, Databases, and Consistency Groups

Setting	Description
Enable File Catalog	When enabled, the Catalog will scan and index the captured data across a Consistency Group.
File Catalog Username and File Catalog Password	When needed, the credentials provided in these spaces grant access to the application being scanned and indexed by the Catalog across a Consistency Group.
Do Not Unmap	<p>Specifies if you want temporary staging disks mapped to the host and used during data movement for backup to remain mapped to the host. LUNs are mapped during the first job and all the subsequent jobs reuse the same mapped LUN.</p> <ul style="list-style-type: none">Keep staging disks mapped between jobs: Select this if you want temporary staging disks mapped to the host and used during data movement to remain mapped to the host. LUNs are mapped during the first job and all subsequent jobs reuse the same mapped LUN. By default, this option is selected. <p>Note: For applications managed using the Actifio Connector where the application is on an OS running inside a VMware VM, this option is ignored. The staging disk is unmapped from the VM after every job.</p> <ul style="list-style-type: none">Unmap staging disks after each job: This option both unmounts the staging disk from the operating system at the conclusion of every job (removing mount points or drive letters), and also unmaps it from the host altogether. This option will require the host to perform a scan for SCSI LUNs at the start of the next job, as the re-mapped staging disks must be rediscovered before they can be remounted.
Truncate Log After Backup	Specify whether to truncate the logs after every backup. When this is enabled, application-related logs are truncated until the recent or current backup. If you truncate logs, you must also back up the transaction log to enable a roll forward recovery.

Table 2: Policy Settings Overrides Valid for SQL Server Instances, AAG, Databases, and Consistency Groups

Setting	Description
Skip Offline Applications in the Consistency Group (For Consistency Group management only)	Specify whether to ignore unavailable applications that are part of a consistency group. You create a consistency group to back up the data of all member applications together to preserve consistency of data across the member applications. Consistency groups are collections of discovered applications from the same host. Options are: <ul style="list-style-type: none"> Fail backup when offline applications are found Skip offline applications during backup
Map staging disks to all Nodes in an Application Cluster (or all ESX Hosts in a Cluster)	If your nodes are in an application cluster, you can use this to ensure that the nodes of an application cluster are protected in case of failover during backup. <ul style="list-style-type: none"> Do not map staging disk to all nodes of application cluster Map staging disk to all nodes of application cluster. In the event of an application cluster failure, this option will protect failover copies.
Backup SQL Server User Logins	Captures the SQL Server instance login records for accounts granted access to databases being backed up. When the database is mounted as a virtual application (application aware mount) the backed up user logins can be optionally restored into the target SQL Server instance, ensuring the virtual database will be accessible by the same users with access to the original source database. Options are Yes or No .
Enable Database Log Backup	The Enable Database Log Backup option allows the SLA policy to backup an Oracle or Microsoft SQL Server database and all associated transaction log files. The logs are backed up when the log snapshot job runs. Options are Yes or No. When set to Yes, the related options are enabled. <p>Note: For details on Log Protection, see Database Log Protection in an SLA Policy on page 23.</p>
RPO	When Enable Database Log Backup is set to Yes , RPO defines the frequency for database log backup. Frequency is set in minutes and must not exceed the database backup interval.
Log Backup Retention Period	When Enable Database Log Backup is set to Yes , log retention is defined separately from the retention of the Snapshot policy. Having a separate retention period allows you to use logs in conjunction with copies of the database stored in both the Snapshot and Dedup pools. The log retention period is a mandatory setting.

Table 2: Policy Settings Overrides Valid for SQL Server Instances, AAG, Databases, and Consistency Groups

Setting	Description
<p>Replicate Logs (Uses StreamSnap Technology)</p>	<p>When Enable Database Log Backup is set to Enable, the Replicate Logs advanced setting allows Oracle archive logs or Microsoft SQL Server database transaction logs to be replicated to a remote VDP appliance. For a log replication job to run, there must be a replication policy (StreamSnap, Dedup-Async, or Remote Dedup) included in the template along with a resource profile that specifies a remote VDP appliance, and at least one successful replication of the database must first be completed. You can then use the logs at the remote site for any database image within the retention range of the replicated logs. This function is enabled by default.</p> <p>Log replication uses StreamSnap technology to perform the replication between the local and remote VDP appliances; log replication goes directly from the local snapshot pool to the snapshot pool on the remote appliance.</p> <hr/> <p>Note: <i>Log replication does not occur until an SQL Server database has been protected and the database has been replicated to the remote VDP appliance.</i></p> <hr/>
<p>Log Staging Disk Growth Size</p>	<p>When Enable Database Log Backup is set to Yes, Log Staging Disk Growth Size defines the growth to use when automatically growing the staging disk on which the logs reside. This setting is from 5 to 100 percent.</p>
<p>Estimated Change Rate</p>	<p>When Enable Database Log Backup is set to Yes, this setting defines the daily change (in percent), which allows the VDP appliance to better calculate the size of the staging disk needed to hold logs. This setting is from 0 to 100.</p>
<p>Compress Database Log Backup</p>	<p>When Enable Database Log Backup is set to Yes, this setting instructs the source database to compress its logs before they are captured by AGM. The database server performs log compression during log backup. Options are Yes or No. When set to Yes, the Compress Database Log Backup option is enabled.</p>
<p>Enforced Retention</p>	<p>Allows the user to configure the desired immutability period between 0 and 36525 days. By default, the value is set to 0 for all existing policies.</p> <p>You can modify a policy that is already used to protect an application by setting a longer enforced retention period. However, you cannot shorten the enforced retention period.</p> <p>You cannot set enforced retention for a StreamSnap policy whose retention is "Only keep the most recent remote image".</p> <p>When configured to send OnVault data to an object store with Enforced Retention integration, images will also be protected against direct deletion by an object storage administrator until they reach the specified enforced retention period.</p> <hr/> <p>Note: <i>Enforced Retention cannot be overridden on a per-application basis. The option does not appear on the "Policy Overrides" page.</i></p> <hr/>

Table 2: Policy Settings Overrides Valid for SQL Server Instances, AAG, Databases, and Consistency Groups

Setting	Description
Job Behavior When Target VM Needs Snapshot Consolidation	<p>Select an action if the VM requires consolidation:</p> <ul style="list-style-type: none"> Fail the job if VM needs consolidation: Point-in-time/DAR/direct-dedup jobs fail. Run the job without performing consolidation: All jobs run normally even if consolidation is pending. Perform consolidation at the beginning of the job: Point-in-time/direct-dedup/DAR jobs try to perform consolidation at the beginning of the job. If consolidation fails, the job fails with an error message.
Fail On Missing Start Path	<p>If one or more start paths are specified, and any of these start paths does not exist, the job will fail with the message <i>UDSAgent: Specified start path does not exist</i>. If no start paths are specified, this option has no effect. Options are Yes or No.</p> <hr/> <p>Note: <i>The default state for this is No (disabled), which is the same behavior of the previous versions of the VDP Connector; the job will not fail if a start path does not exist.</i></p>
Enable Degraded Capture Mode	<p>Degraded capture mode captures incremental data when Change Block Tracking (CBT) service is unavailable. Data capture may take longer. Options are Yes or No.</p>
Script Timeout	<p>The VDP Connector allows you to create host-side scripts that run on an application's host before and/or after a policy is run. The four timeouts provided in a policy template map directly into the four stages of a host-side script.</p> <hr/> <p>Note: <i>By default, the script timeout values are 1. If a script timeout is not specified, the value will be blank.</i></p> <hr/> <p>Script Init Timeout: Defines how long a policy should wait before assuming host-side scripts on a managed host have been initialized. 120 seconds is the default value, allowed range is from 1 to 86400 seconds (24 hours).</p> <p>Script Freeze Timeout: Defines how long a policy should wait before assuming the application is frozen and ready for data capture. 60 seconds is the default value, allowed range is from 1 to 86400 seconds.</p> <p>Script Unfreeze Timeout: Defines how long a policy should wait before assuming the application is unfrozen. 60 seconds is the default value, allowed range is from 1 to 86400 seconds.</p> <p>Script Finish Timeout: Defines how long a policy should wait before data capture is complete. 60 seconds is the default value, allowed range is from 1 to 86400 seconds.</p> <p>Script Post Replication Timeout: Defines how long a policy should wait before replication is complete. 60 seconds is the default value, allowed range is from 1 to 86400 seconds.</p>

5 Mounting a Microsoft SQL Server Database

This chapter describes how to mount captured Microsoft SQL Server data in three ways:

- Instances
- Primary database of an AlwaysOn Availability Group
- Consistency groups
- Individual members of a consistency group
- System databases
- User databases

This chapter includes:

[Mounting Captured Microsoft SQL Data](#) on page 29

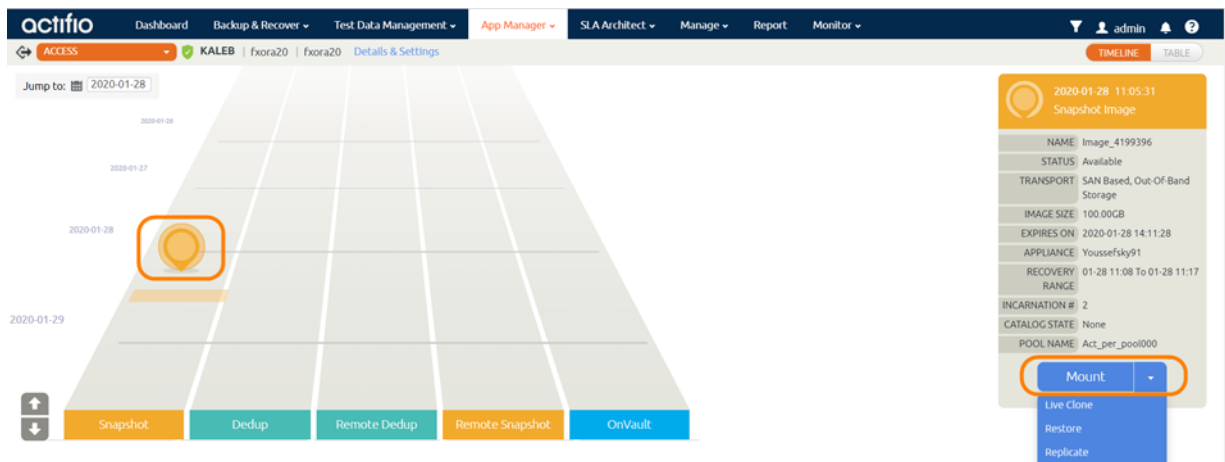
[Mounting an SQL Server Database as a New Virtual Database](#) on page 31

[Mounting Encrypted SQL Data](#) on page 34

Mounting Captured Microsoft SQL Data

With a standard mount you can mount the Microsoft SQL data to another server where it can be picked up and used by another Microsoft SQL Server database. To mount just the captured Microsoft SQL data:

1. Open the App Manager to the Applications list.
2. Right-click an SQL instance, user database, system database, cluster, or availability group and select **Access**. The filters in the left-hand pane make it easier to find the database you need.
3. On the runway, select the specific image to be mounted. On the right side, select **Mount**.




Selecting an SQL Image to Mount

Note: You can use the calendar widget in the upper left corner to narrow the range of backup images.

4. On the Mount page, fill in the **Details & Settings** at the top of the page as needed.

Note: OnVault images can be mounted very efficiently from Actifio Appliances. Once the entire image is copied, then the mount is performed from the snapshot pool.



Initial Mount Screen for an SQL Server Instance

5. Enter a label that will allow you to clearly identify this mounted data.
6. In the **Application Options** section of the Mount Image dialog box, do **NOT** select **Create New Virtual Application**. (To mount the Microsoft SQL data as a virtual database, see [Mounting an SQL Server Database as a New Virtual Database](#) on page 31.)
7. Fill in **Mapping Options** and **Item Selection** as needed for this new database.
Options presented vary according to the source that is selected. For example, databases on VMware VMs will have a **Map to all ESX hosts** option. Clustered databases will have the **Map to Cluster Nodes** option.
8. Click **Submit** and the mount job is submitted.
9. Once the mount operation is successful, log on to the database server and verify that the mounted image is available.

Mounting an SQL Server Database as a New Virtual Database

A Virtual Application Mount operation mounts a captured application as a virtual application. It allows you to bring a database online quickly without having to actually move the data and without having to manually configure a new instance of the database. A Virtual Application Mount addresses the challenges of creating and managing copies of production databases without manual intervention by database, server, and storage administrators.

This chapter describes how to mount a captured Microsoft SQL Server database as a virtual application. You can mount Microsoft SQL:

- Instances
- System databases
- User databases
- Consistency groups
- Individual members of a consistency group
- Primary database of an AlwaysOn Availability Group

Note: Before mounting an image, ensure that the WWPN/iSCSI port of the host where the images will be mounted is accessible to the Actifio Appliance.

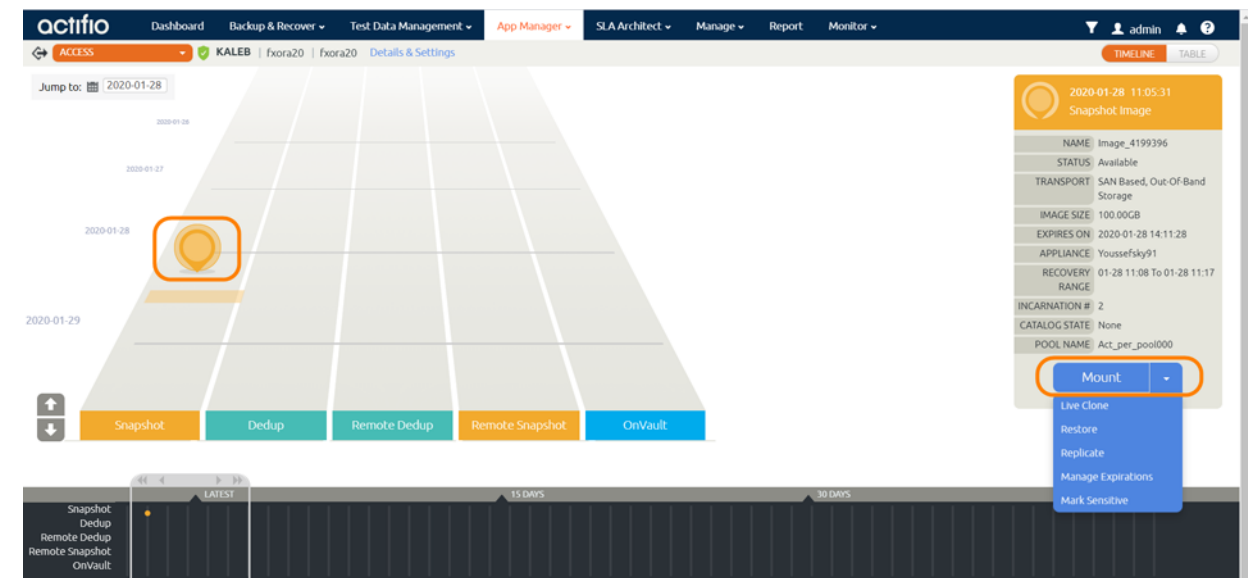
Note: A database mounted to a protected SQL Instance as a virtual application will not be protected with the SQL Instance.

If you want to mount just the Microsoft SQL data for database recovery, see [Mounting Captured Microsoft SQL Data](#) on page 29.

For corrupt or deleted databases, mounting an SQL database as a virtual application to its original server is an efficient alternative to performing a restore of the database.

To mount a captured Microsoft SQL Server database as a virtual application:

1. Open AGM to the **App Manager > Applications** list.
2. Right-click an SQL instance, user database, system database, cluster, or availability group and select **Access**. The filters in the left-hand pane make it easier to find the database you need.
3. On the runway of images, select the image to be mounted. On the right side, select **Mount**.




Note: You can use the calendar widget in the upper left corner to narrow the range of backup images.

4. On the Mount page, fill in the **Details & Settings** as needed for this database.

Note: OnVault images can be mounted very efficiently from Actifio Appliances. Once the entire image is copied, then the mount is performed from the snapshot pool.

5. In the **Application Options** section, enter a label that will allow you to clearly identify this mounted data.
6. Select **Create New Virtual Application**.
7. Set the **Mapping Options** and **Mount Location** information as required for this new database.



Filling In the Mount Options

8. If the database was captured along with its logs, the App Options dialog box provides an option to roll the logs to a specific point in time based on when and how often logs were captured.
9. From the **SQL Server Instance Name** drop down list, select the SQL Server instance that will manage the new virtual application. If the required instance name is not included in the drop down, you can manually type the name in the space provided.
10. From the **SQL Server Database Name** drop down list, specify a name for the new SQL Server database. Valid characters include letters, numbers, @, #, -, _ . Leading and trailing spaces are not allowed.
11. The Virtual Application Mount will be a new virtual database. To protect the new virtual database, select **Manage New Application** and select the template and profile to use.

The Virtual Application Mount will be a new database. The snapshots of the database are incremental unless you apply a policy template with Force Out-of-Band Backup checked.

Note: There is one exception to this: if the target server is a VMware VM, you must select "pRDM" when performing the mount if you want the child database to have the efficient incremental snapshots. If you leave the default of "vRDM", then the first snapshot job will be a full backup.

If you mount the virtual application to a host known to the Actifio Appliance, then the virtual application will appear in the Application Manager list of applications.

If you do not select **Manage New Application** then it will appear in the Application Manager as an unprotected application. It can be protected like any other application. In this case, the first snapshot will be a full image.

Virtual SQL databases mounted to an SQL instance must be protected separately from the instances user and system databases.

12. **Recover Database After Restore:** Leave this selected (default) if you want to bring the database to an online state ready to process transactions. Deselect this to leave the database in a restoring state, so additional transaction log backups can be applied to roll the database forward to a specific point in time.
13. **Recover User Logins:** If you have enabled the policy option "Backup SQL Server Logins", then selecting this option will result in a restore of those logins to the target SQL Instance. For domain accounts, the user accounts will only be restored if the target instance is on a server in the same domain, or in a domain with a trust relationship to the source SQL Server. SQL Local accounts will always be restored. Use this option if you want to ensure all users who could access the original source database can access the new virtual database.
14. In the **Username and Password** fields enter a user name and password as needed. If the Actifio Connector account does not have privileges to detach the database during an unmount operation or to apply transaction logs, then enter credentials here for an account with those privileges. See [Chapter 2, Required SQL Server Roles for the Windows User](#) for details.
15. Select a **Recovery Model** which can be the same as the source database or different.
16. For **Overwrite Existing Database**, indicate when to overwrite a database on the target server that has the same name as the new database(s) being mounted: Yes, No, or Only if it's Stale.
17. In the Mapping Options, you can enter a **Mount Location**. If an application has only one volume, you must specify the mount location here. If an application has multiple volumes, you can:
 - o Enter a mount location: all volumes will be automatically mounted at the specified mount location.
 - o Leave this space blank and in the **Advanced Options**, manually specify a mount location for each volume.

Note: Select all volumes for the application. Data on all volumes is needed to mount the database correctly.

18. Click **Submit** and the job is submitted.
19. When the mount job is finished, log onto the database server and verify that mounted image is available.

Note: When performing a Virtual Application Mount of a SQL Server database to a SQL Server Failover Instance, if you specify a custom mount point, the custom mount point must reside on a volume that is a cluster resource. This is required to allow the SQL Server Instance to move to other cluster nodes in case of failover.

Mounting Encrypted SQL Data

Actifio Appliances capture encrypted SQL Server databases but do not capture their private keys, encryption certificates, or passwords.

This chapter describes:

- [Determining if SQL TDE is Enabled](#) on page 34
- [Troubleshooting SQL Server Encryption](#) on page 36
- [SQL Server Master Key, Encryption Certificate, and Password Procedures](#) on page 37

If you are restoring an encrypted SQL Server database over an existing SQL Server database, the private key, encryption certificate, and password are already present on the SQL Instance and once the restore operation finishes, the SQL Server database will work as expected.

If you are performing a Virtual Application Mount of an encrypted database, or a mount of just the encrypted SQL data, the SQL instance on which the encrypted database or data will be mounted must have:

- Transparent Data Encryption (TDE) enabled
- A copy of the Private Key from the source SQL Server database
- A copy of the encryption certificate from the source SQL Server database
- Provide the password of the source SQL Server database

Procedures are in [SQL Server Master Key, Encryption Certificate, and Password Procedures](#) on page 37.

Note: *If you are not mounting the database back to the source SQL instance, then the private key and encryption certificate must be manually copied from the source SQL instance to the new SQL instance.*

Determining if SQL TDE is Enabled


To determine if TDE is enabled on an SQL instance, you can use Microsoft's SQL Server Management Studio's user interface ([SSMS](#)), or you can use a manual query to determine if encryption is enabled on a database. For example:

```
SELECT
  DB_NAME(database_id)AS dbname,
  encryption_state,
  case encryption_state
    WHEN 0 THEN 'Unencrypted (no database encryption key present)'
    WHEN 1 THEN 'Unencrypted'
    WHEN 2 THEN 'Encryption in Progress'
    WHEN 3 THEN 'Encrypted'
    WHEN 4 THEN 'Key Change in Progress'
    WHEN 5 THEN 'Decryption in Progress'
    ELSE CAST(encryption_state AS varchar(20))
  END AS encryption_state,
  key_algorithm,
  key_length
FROM sys.dm_database_encryption_keys
```

SSMS

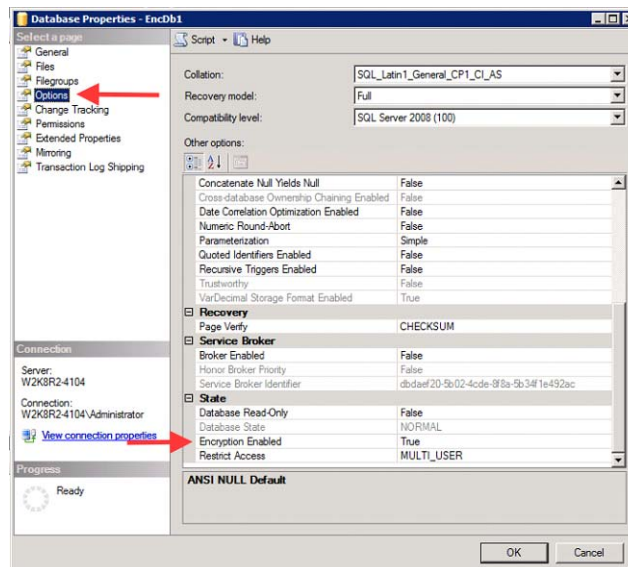
To use SSMS to determine if encryption is enabled on a database:

1. From SSMS right click on the database name:



Selecting Properties

2. From the drop down menu select **Properties** and the database's properties are displayed:



Database Properties

3. On the left-hand side of the Properties page, under **Select a Page**, click **Options** and the options for the database are displayed.
4. Under **State**, ensure that **Encryption Enabled** is set to **True**.

Troubleshooting SQL Server Encryption


Two common errors are:

SQL error 24583: missing encryption certificate

SQL error 33117: Transparent Data Encryption not enabled

SQL error 24583: missing encryption certificate

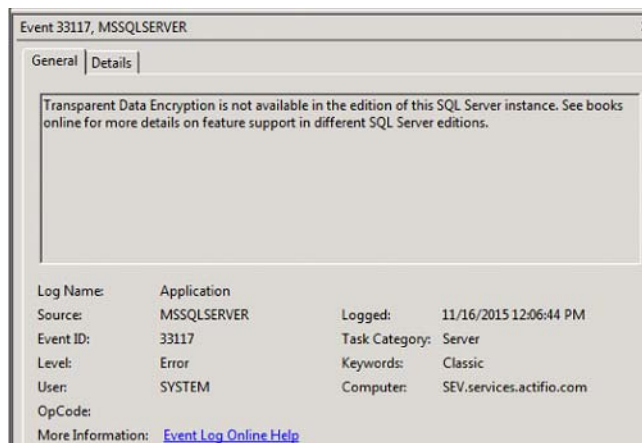
The following 24583 SQL error indicates that you are trying to perform a mount to an SQL instance that does not have the encryption certificate of the source SQL instance:



Example Error

SQL error 33117: Transparent Data Encryption not enabled

The following 33117 SQL error indicates that you are trying to perform a mount of an encrypted SQL Server database to an SQL instance that does not have Transparent Data Encryption enabled:



Example Error

SQL Server Master Key, Encryption Certificate, and Password Procedures

Creating and copying master keys and encryption certificates are standard Microsoft SQL procedures that are not unique to Actifio Appliances. They are provided here as a convenience:

[Create a New Master Key](#)

[Create a New Encryption Certificate](#)

[Apply Server Master Key and Encryption Certificate](#)

[Create Copy of Server Master Key, Encryption Certificate, and Provide Source Password](#)

[Copy Encryption Certificate, Private Key, and Provide Source Password](#)

For more information, see Microsoft's detailed information on security certificates and keys:
<https://msdn.microsoft.com/en-us/library/ff848768.aspx>

Create a New Master Key

```
use master;
go
create master key encryption by password = 'SMKSourcePassword';
go
```

Create a New Encryption Certificate

```
use master;
go
create certificate sourcedbcert with subject = 'Act Test Cert';
go
```

Apply Server Master Key and Encryption Certificate

```
use DATABASENAME;
go
CREATE DATABASE ENCRYPTION KEY
WITH ALGORITHM = AES_128
ENCRYPTION BY SERVER CERTIFICATE sourcedbcert;
go
alter database DATABASENAME
set encryption on;
go
```

Create Copy of Server Master Key, Encryption Certificate, and Provide Source Password

If an SQL Server database on one SQL instance will be mounted to another SQL instance, you must manually copy the to-be-mounted database's Server Master Key, encryption certificate, and password. Then copy the Server Master Key, encryption certificate, and password to the other SQL instance. To make a copy of a Server Master Key, Encryption Certification, and password:

```
use master;
go
backup certificate sourcedbcert to file = 'E:\Enc\Sourcecert'
with PRIVATE KEY (file='E:\Enc\Privatekey',
ENCRYPTION BY PASSWORD='SecurePassword');
go
```

Copy Encryption Certificate, Private Key, and Provide Source Password

If an encryption-enabled SQL Server database or data will be mounted to a new SQL instance, the new instance must have a copy of the source SQL instance's Server Master Key, encryption certificate, and password. Manually copy the encryption certificate and password copies you made on the source SQL instance in the previous section to the new SQL instance. From the new SQL instance:

```
create certificate destinationdbcert
FROM file = 'C:\Program Files\Actifio\sqlenc\Sourcecert'
with private key (file = 'C:\Program Files\Actifio\sqlenc\Privatekey',
decryption by password = 'SecurePassword')
go
```

6 Mounting Databases into SQL AlwaysOn Availability Groups

This chapter details:

- [Creating an SQL Server AAG in an Actifio Snapshot Pool](#) on page 39
- [Creating an SQL Server AAG Outside of An Actifio Snapshot Pool](#) on page 39
- [Creating the New SQL Server AlwaysOn Availability Group](#) on page 40

Creating an SQL Server AAG in an Actifio Snapshot Pool

This approach is the fastest way to create a new SQL Server AlwaysOn Availability Group for short-term use. Because the mounted application will run in the Actifio Snapshot Pool, it will achieve the same performance as data captured in the Snapshot Pool.

Caution! *Ensure that there is enough space in the Snapshot Pool to accommodate the AAG and regular snapshots.*

1. One at a time, mount each member of the AAG as a new virtual database. See [Mounting an SQL Server Database as a New Virtual Database](#) on page 31.
2. Select the required host.
3. Select the SQL instance.
4. Enter a database name. Use the same name for each mounted image.
5. Repeat the process for each AAG member.
6. Once each of the AAG members are mounted as virtual databases, via SQL select the SQL instance that will be the primary database and recover it:

```
recover database <name> with recovery
```
7. Keep the secondary database copies in “restoring” state.
8. When the primary database has been restored, via SQL, create the new AAG and join the primary and secondary databases as described in [Creating the New SQL Server AlwaysOn Availability Group](#) on page 40.


Creating an SQL Server AAG Outside of An Actifio Snapshot Pool

To create an AAG that resides outside of an Actifio Snapshot pool, as described in [Chapter 8, Cloning SQL Server Databases](#).

Creating the New SQL Server AlwaysOn Availability Group

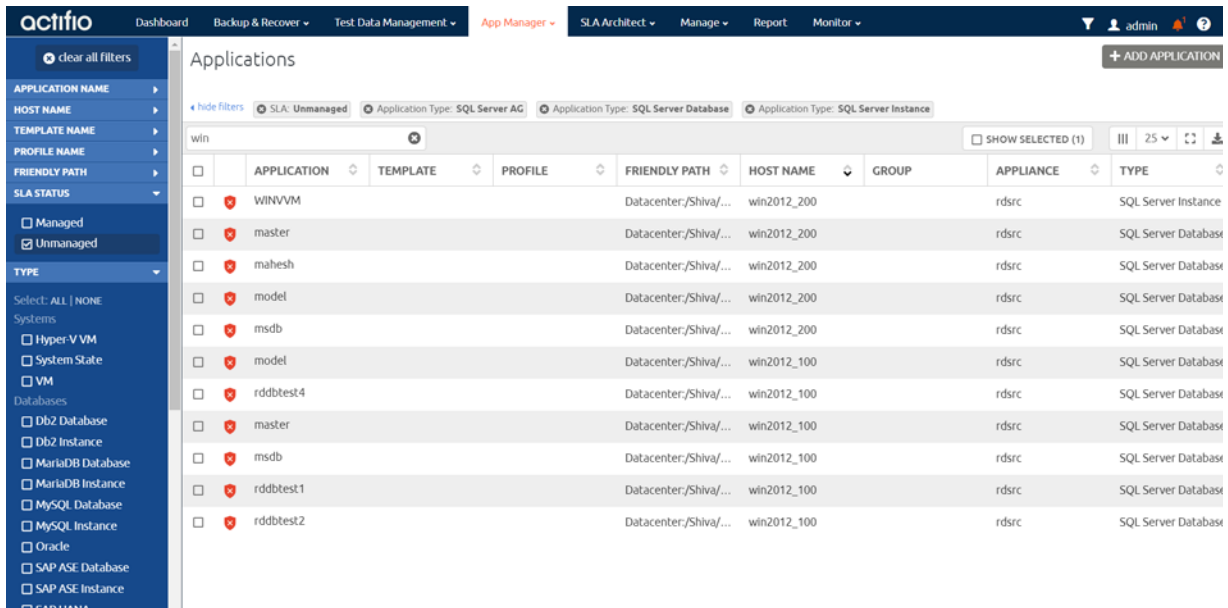
After all databases for the new AAG have been recovered, use SQL Studio or T-SQL to create the new AAG:

1. From SQL Studio create a new AAG.



Create New AAG in SQL

2. Add the primary and secondary databases to the AAG.
3. Click **Next**. You will be prompted to select a synchronization method. Select **Join only**: The Join only operation will form the primary and secondary databases into an AAG.
4. Monitor the progress of the operation. When the operation is complete, both the primary and secondary databases are in sync.
5. If it is not already present, install the appropriate Actifio Connector on AAG's host.
6. If auto discover applications is not enabled, manually discover applications on the AAG host.
7. After discovery is complete, the host's AAGs will appear in the Application Manager:



Discovered AAGs

8. Capture members of an AAG just like any other database. See [Chapter 4, Capturing Microsoft SQL Server Instances and Databases](#) for details.

7 Mounting and Migrating SQL Data

The SQL Mount and Migrate feature is available for recovery of:

- SQL databases (stand-alone and failover cluster)
- Consistency Group of SQL Databases
- SQL Instances
- SQL AGs

A Mount and Migrate operation allows you to restore an application with near-zero downtime by first mounting it locally, and then migrating it to the original location or to a new location. Users have normal access to the application while it is mounted, and the migration step is very fast. This is similar to a VMware Storage VMotion operation.

There are two common use cases that are made much faster by Actifio Mount and Migrate:

- **Database Recovery:** You can recover SQL Server and file system data instantly using Actifio's existing capabilities, and then migrate the data in real-time to production storage, while the database is up and running.
- **Creating a Database Copy on Other Storage:** You can move recovered data into other local or SAN storage while the databases are up and running, completing the process with almost no downtime.

The three-step migration process

SQL Mount and Migrate is a three step process that includes:

- **Step 1: Mount or Restore:** Perform either a Virtual Application Mount or a mount and migrate restore.
- **Step 2: Scheduling the Migration:** Configure the migration schedule for an image from the Active Mounts page. Image migration follows the schedule defined and migrations jobs are run repeatedly at the frequency specified.
- **Step 3: Finalize:** Initiate the last and final stage of the image migration process.

Note: When working with SQL AAG, the finalize step cannot be performed on a database in a group that is on the primary node. To finalize migration on the primary node, a failover to another node must be performed first, which changes the node to be a secondary.

Step 1: Mount or Restore

You can perform either a Virtual Application Mount or a Mount and Migrate restore on the desired image. For instructions on how to perform a Virtual Application Mount of SQL databases, see [Mounting an SQL Server Database as a New Virtual Database](#) on page 31.


For instructions on performing a Mount and Migrate restore, continue to [Mount and Migrate Restore](#), below.

Multiple restores can be performed concurrently for a parent application like SQL Instance, SQL AAG, or SQL CG, but they have to be on different databases. When more than one Mount and Migrate jobs are in progress, you will see a notification in a yellow banner.

Mount and Migrate Restore

To perform a Mount and Migrate restore:

1. Click the App Manager tab and select **Applications** from the drop-down menu. The Applications page opens.
2. Right-click the SQL application with the image that you want to restore using Mount and Migrate and select **Access**. The Access page opens listing captured images.
3. Click the image, then select **Restore** from the list of operations on the right side of the page. The Restore page opens.
4. Select the **Mount and Migrate** option. Depending on your database, you may see more or fewer options than what you see in the image below.



Restore

Use this page to initiate a restore operation. A restore will take the existing database offline and overwrite their data files.

TRADITIONAL MOUNT AND MIGRATE

5. If the selected database does not have logs, the Restore page does not show roll forward options. If the SQL Server database was managed with a Log Protection SLA template, and logs are available with the image and you want to use them to roll forward to a specific point in time, you can:
 - o Specify to roll forward using either User Time or Host Time. You can base the dates and times on User Time or Host time. User Time is relative to the viewer of the current screen. Host time is relative to the system that hosts the data to be restored.
 - o Use the Calendar tool to select a date from which to initiate the restore operation.
 - o Use the Restore Range slider to select a specific point in time to restore the database. Slide the slider tool all the way to the left to restore just the SQL Server database.

6. In Label, optionally enter a name. The Label is pre-populated with the text "Restore - M&M - <current date and time>", so if you do not update the property, the default value will be used.
You will not be able to change the default selection for Restore with Recovery.
7. For SQL Server Instance Name, select the target SQL Server Instance. The new database will be managed by the Instance you specify. Similarly, for Consistency Groups, enter the name of the target Consistency Group in the Enter Consistency Group Name field.
8. In SQL Server Database Name, enter the new SQL Server database to be provisioned.
9. For SQL Instances and Consistency Groups, select the databases to be included in the restore job. The Restore with Recovery option is enabled by default and you cannot disable it.
10. In Username, enter the name of the username for database provisioning. This is needed only when the account running the Connector does not have the privileges to apply transaction logs or to dispatch a database.
11. In Password, enter the password for the user you specified in [Step 10](#).
12. Click **Submit**. A warning dialog opens. Read it and then enter DATA LOSS to confirm.

The selected databases are taken offline. A mount is performed to provide fast access to the databases with the selected point in time. After the mount completes, the "migrate" option becomes available as an action on the resulting active image.

Note: *If the target is a VM, and the mount mode is not pRDM, the user will see the warning text below the "Data Loss" message: "Warning: pRDM is not selected. This means that backups taken while the database(s) are running from the mount will be captured out-of-band and will consume a full copy's worth of space in the snapshot pool, plus changes. When the migrate begins, backups will stop until the migration is finalized."*

Step 2: Scheduling the Migration

You can configure the migration schedule from the Active Mounts page. Any Virtual Application Mounted image can be migrated at a later time. Images that were restored using the Mount and Migrate option can also be migrated. Mounted images show the Image State of:

- Mounted images show image state Mounted
- Images that were restored using Mount and Migrate have image state Restore (Mounted)

Once migration starts, the image state changes to Migrating (if the image is virtual database mount) or Restore (Migrating) (if the image was created using the Restore Mount and Migration option). The image remains in Migrating state until you initiate finalize migration in step 3, when the last migration is performed.

Select whether you are:

[Configuring the Migration Schedule for Mount \(Restore\) Images](#) on page 44

[Configuring the Migration Schedule for Mounted Images](#) on page 45


[Updating Migration Frequency](#) on page 46

[Canceling a Migration Job](#) on page 46

Configuring the Migration Schedule for Mount (Restore) Images

To configure the migration schedule:

1. Select an active image with Image State of Mounted or Restore Mounted.
2. Select **Migrate** from the drop-down menu at the bottom of the page. The Update Migrate Frequency page opens.



Update Migration Frequency

The form contains two fields: 'FREQUENCY' with a slider set to 9h, and 'COPY THREAD COUNT' with a text input set to 15. There are 'Cancel' and 'Submit' buttons at the bottom right.

3. For Frequency, use the slider to define the frequency with which to schedule migrate jobs from one hour to 24 hours.
4. For Copy Thread Count, specify the number of parallel copy threads to use, per disk volume, during the migration. The default value is four(4).
5. Click **Submit**. You will be prompted to confirm.
6. Click **Proceed** to apply the configuration.
7. You will see a success message after configuration is complete.
8. Click **Okay** in the success dialog to the Active Mounts page.

Configuring the Migration Schedule for Mounted Images

To configure migration schedule:

1. Select an active image with Image State of Mounted or Restore Mounted.
2. Select Migrate from the drop-down menu at the bottom of the page. The Update Migrate Frequency page opens. You may see more options depending on the specific database:

Metadata Card:

2020-02-04 03:19:15	mapped Image
NAME	Image_0594990
APPLICATION	AG_SQL12_Grp1
MOUNTED	Sql2012clust21.Sqa.Actifio
HOST	Com
APPLIANCE	Skysr6
IMAGE STATE	Mapped
CHILD APPLICATION	DB22_mnt22_2

Migrate Configuration:

FREQUENCY: Slider set to 24h.

RENAME FILES TO MATCH NEW DATABASE NAME:

COPY THREAD COUNT:

File Locations:

- Copy files to the same drive/path as they were on the source server
- Choose new file locations at the volume level
- Choose new locations at the file level.

Volumes Table:

SOURCE VOLUME	TARGET VOLUME
F:\	F:\
C:\	C:\

Buttons: Cancel, Submit

3. For Frequency, use the slider to define the frequency with which to schedule migrate jobs (in hours). Minimum value is one (1) hour and the maximum value is 24 hours.
4. For Copy Thread Count, specify the number of parallel copy threads to use, per disk volume, during the migration. The default value is four(4).
5. In the File Locations section, specify whether you want to copy the database files to the same path as the source server, or to a new location. There are three options:
 - o Copy files to the same drive/path as they were on the source server (default option).
 - o Choose new file locations at the volume level.
 - o Choose new file locations at the file level.

If you select New File Locations at the Volume Level, you see a table with the source volume and target volume drop-down.

File Locations:

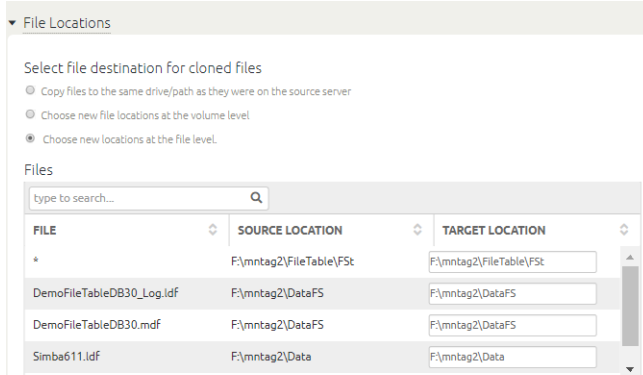
- Choose new file locations at the volume level.

Volumes Table:

SOURCE VOLUME	TARGET VOLUME
F:\mntag2\	F:\mntag2\

In Target Volume, select a target volumes from the drop-down list of all discovered file system applications. If needed, you can also type the volume, for example: M:\, or L:\Logs\Log1.

If you select **New File Locations at the File Level**, you see a table with three columns: **File**, **Source Location**, and **Target Location**.



6. In **Target Location**, enter the new file location and make other modifications as needed.
7. Click **Submit**. You will be prompted to confirm.
8. Click **Proceed** to apply the configuration.
9. You will see a success message after configuration is complete.
10. Click **Okay** in the success dialog to the **Active Mounts** page.

Once migration starts, the image state changes to **Migrating** (if the image is virtual database mount) or **Restore (Migrating)** (if it was created using the **Restore Mount and Migration** option). The image will be in **Migrating** state until you initiated **finalize migration** in step 3, when the last migration that is performed.

Updating Migration Frequency

You can go back and update an existing migration frequency if needed. To update the migration schedule:

1. Select an active image with Image State of **Mounted** or **Restore Mounted**.
2. Select **Migrate** from the drop-down menu at the bottom of the page. The **Update Migrate Frequency** page opens.
3. For **Frequency**, use the slider to define the frequency with which to schedule migrate jobs (in hours). Minimum value is one (1) hour and the maximum value is 24 hours.
4. For **Copy Thread Count**, specify the number of parallel copy threads to use, per disk volume, during the migration. The default value is four(4).
5. For **Mounted** images, update the **File Locations** section using instructions in **Configuring Migration Schedule for Mount and Migrate**.
6. Click **Submit**. You will see a success message.
7. Click **Okay** to close the message and return to the **Active Mounts** page.

Canceling a Migration Job

You can cancel migration any time before you initiate the **Finalize Migration** process. To cancel an image migration:

1. Select an active image with Image State of **Migrating**.
2. Select **Cancel Migration** from the drop-down menu at the bottom of the page. You will see a warning message.
3. Click **Proceed** to cancel. This stops the migration and deletes all data copied over during previous migrations.
4. Go to **Monitor > Jobs**, and filter by **Cancel (Migrate)** job type if you want to view the job details.

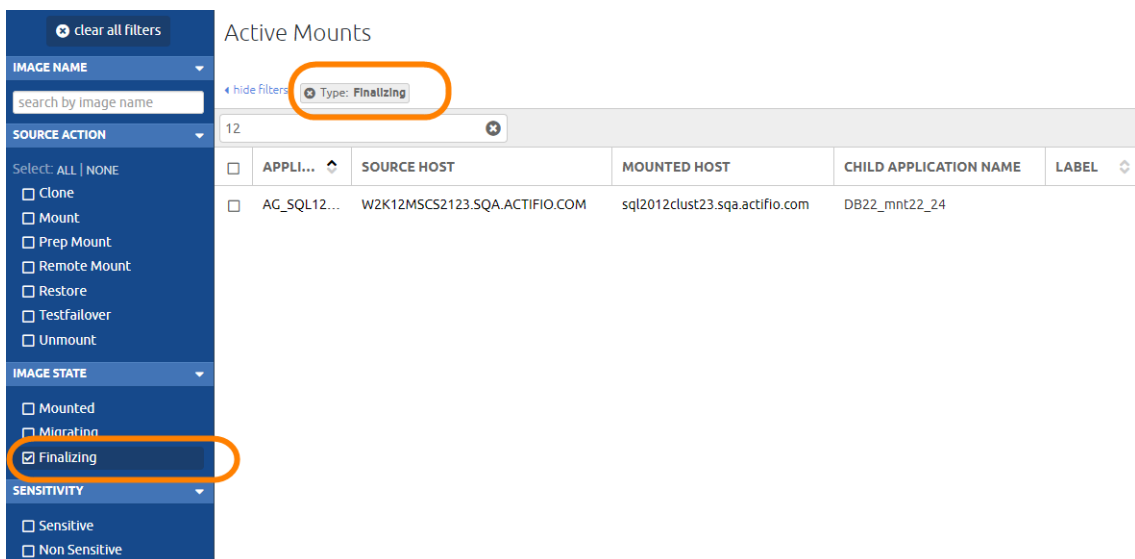
Step 3: Finalize

Finalize Migration initiates the last and final stage of the image migration process. Once you begin finalizing, you cannot cancel the migration process.

To begin the Finalize step:

1. Go to the Active Mounts page and filter by Image State of Migrating or Restore (Migrating).
2. Select the desired image and then select **Finalize Restore** from the drop-down menu at the bottom of the page.
3. Read the confirmation message and then click **Proceed**. The databases are taken offline during final migration and then brought back up again on the target production database.
4. To view progress of the job, go to Monitor > Jobs, and filter by Finalize job type. Locate the job and view the job details.

You can also view images with the Finalize Image State in the Active Mounts page:



The screenshot shows the 'Active Mounts' interface. On the left, there is a filter sidebar with sections for 'IMAGE NAME', 'SOURCE ACTION', 'IMAGE STATE', and 'SENSITIVITY'. The 'IMAGE STATE' section has 'Finalizing' selected and circled in orange. The 'SOURCE ACTION' section has 'Finalize Restore' selected and circled in orange. The main table displays one row of data with columns: APPLI..., SOURCE HOST, MOUNTED HOST, CHILD APPLICATION NAME, and LABEL. The row contains: AG_SQL12..., W2K12MSCS2123.SQA.ACTIFIO.COM, sql2012clust23.sqa.actifio.com, and DB22_mnt22_24.

APPLI...	SOURCE HOST	MOUNTED HOST	CHILD APPLICATION NAME	LABEL
AG_SQL12...	W2K12MSCS2123.SQA.ACTIFIO.COM	sql2012clust23.sqa.actifio.com	DB22_mnt22_24	

8 Cloning SQL Server Databases

You can clone (copy) a captured SQL Server database, instance or AG image to any physical or virtual host managed by your VDP appliance, including virtual cluster hosts (when the option to mount to a virtual SQL cluster host is enabled), physical cluster hosts, vCenter hosts, SCVMM hosts, and Hyper-V hosts). You cannot clone to ESX hosts, IBM HMC hosts, HP-UX hosts, TPGS hosts and OpenVMS hosts.

The cloning process varies slightly depending on whether you are cloning a single database image such as a member of an Always on Availability Group (AAG) or multiple images in an SQL instance.

Use a clone operation:


- If you have copies of multiple SQL Server databases on a single volume, to avoid unintentional data loss as the contents of the entire volume get overwritten during restore of the volume.
- If the original database has been removed because of corruption, or if the old database server is being replaced with a new server.
- If you are restoring databases in a consistency group, all databases in the consistency group are overwritten. If you do not want to overwrite all databases, clone a single database.

To clone a SQL database to a host:

1. Open the App Manager to the **Applications** list.
2. Right-click the database with the image that you want to clone, then choose **Access**. The Access page opens listing captured images in the Timeline ramp view.

The screenshot displays the Actifio App Manager interface. The top navigation bar includes 'actifio', 'Dashboard', 'Backup & Recover', 'Test Data Management', 'App Manager', 'SLA Architect', 'Manage', 'Report', and 'Monitor'. The main content area shows a 'Timeline' view for 'AGM-W2K8SQL2008'. The timeline is a ramp view with columns for 'Snapshot', 'Dedup', 'Remote Dedup', 'Remote Snapshot', and 'OnVault'. A specific image is highlighted with an orange circle, and a context menu is open over it, showing options: 'Mount', 'Live Clone', 'Clone', 'Restore', 'Replicate', 'Manage Expirations', and 'Mark Sensitive'. The 'Clone' option is highlighted with an orange circle. A sidebar on the right shows details for the selected image, including 'NAME: Image_4336274', 'STATUS: Available', 'TRANSPORT: SAN Based, Out-Of-Band Storage', 'IMAGE SIZE: 179.88GB', 'EXPIRES ON: 2020-02-18 20:02:13', 'APPLIANCE: Sky7.1.2.DJ', 'CATALOG STATE: None', and 'POOL NAME: Act_per_pool000'.

3. Select an image and then select **Clone** from the list of access operations. The Clone page opens.



4. Select a target from the Target drop-down list.
5. Enter a unique name for the new clone in the Label field.
6. If necessary, change the storage pool from the Storage Pool drop-down list. The default storage pool is act_per_pool (the Snapshot Pool).
If you are cloning multiple SQL databases into a consistency group, you can append a suffix and/or a prefix to the database's name.
7. Under Application Options, select one or more databases to clone. Unlike the Mount operation, no new Consistency Group is created when multiple databases are cloned.
8. If the SQL server database, instance or AG is managed with a Log Protection SLA template, and logs are available with the image, you will see the Roll Forward Time option. To use the logs to roll forward to a specific point in time, you can:
 - o Specify to roll forward using either User Time or Host Time. You can base the dates and times on User Time or Host Time. User Time is relative to the viewer of the current screen. Host Time is relative to the system that hosts the data to be restored.
 - o Use the Calendar tool to select a date from which to initiate the restore operation.
 - o Use the Restore Range slider to select a specific point in time to restore the image. Slide the slider tool all the way to the left to restore just the SQL AG.

Note: When performing a clone from OnVault, a roll-forward range is displayed only when logs are available on the local appliance. This includes the scenario where an OnVault import was performed on the same appliance used as a target for dedup, DAR, or StreamSnap replication with log replication enabled.

9. In SQL Server Instance Name, select a target SQL Server instance to manage the new database.
10. In SQL Server Database Name, enter a name for the new SQL Server database to be provisioned.
11. Enable the Rename Files to Match New Database option if you want to rename the database files to match the new database name(s).
12. In the Advanced Options section, enter information for the additional fields required. Fields marked with an asterisk (*) are required.

Property	Description
Recover Database After Restore	If Recover Database After Restore is not enabled, the SQL Server database is left in a state where logs can be rolled forward. When it is enabled, the SQL Server database is brought online and logs cannot be rolled forward beyond the time specified in the mount.
Recover User Logins	This applies only if Backup SQL Server User Logins in the Policy Settings or Policy Settings Overrides is set to Yes (this is not the default). If that is set to Yes, all user logins backed up from the source instance will be restored into the target instance. Domain accounts will only restore if the target SQL Server is in the same domain or forest as the source and if any required trust relationships are in place. See Backup SQL Server User Logins on page 25.
User Name/Password	User credentials for database provisioning. The User Name is only required when the account running the Actifio Connector (typically "Local System") does not have privileges to apply transaction logs, or to detach a database (which is typically required during a subsequent unmount).
Overwrite Existing Database	Overwrites the original database.

13. In the File Locations section, specify whether you want to copy the database files to the same path as the source server, or to a new location. There are three options:
 - o Copy files to the same drive/path as they were on the source server (default option).
 - o Choose new file locations at the volume level.
 - o Choose new file locations at the file level.

If you select the second option (new file option at the volume level), you will see a table with the source volume and target volume drop-down.

In Target Volume, select a target volumes from the drop-down list of all discovered file system applications. If needed, you can also type is the volume, for example: M:\, or L:\Logs\Log1.

If you select the third option (new file locations at the file level), you will see a table with three columns: File, Source Location, and Target Location.

14. In Target Location, enter the new file location as needed.
15. Click **Submit**. A job is submitted to clone the image to the selected host. You can verify that the clone operation is successful by viewing the job status in System Monitor. Once the clone job is complete, the image becomes active and is available in the Managing Active Mounts view of the Application Manager.

9 Restoring SQL Server Databases

If a database was deleted or corrupted, you have the option of performing either a full restore operation, creating a clone, or mounting the database almost instantly as a virtual application and then migrate it back to the original location or to a new location. To mount and migrate the database, see [Chapter 7, Mounting and Migrating SQL Data](#).

This chapter describes:

- [Microsoft SQL Server Database Restore Overview](#) on page 54
- [Restoring Microsoft SQL Instances and Databases](#) on page 55
- [Restoring a SQL Server Database to a Different Host](#) on page 56
- [Restoring SQL Server Databases in a Consistency Group](#) on page 56
- [Restoring SQL System Databases](#) on page 57
- [Restoring to an SQL Server Cluster](#) on page 58

Note: Do not use the procedures in this chapter if you have copies of multiple SQL Server databases on a single volume. This may result in unintentional data loss as the contents of the entire volume get overwritten during restore of the volume.

If the original database has been lost, you can mount an image back to the database server, copy files from the backup image to their original location, and then attach the database. Once the database is attached, you can rerun the restore operation if you want to roll forward the database logs.

The restore process is wizard driven and varies slightly depending on whether you are restoring a single database image such as a member of an Always on Availability Group (AAG) or multiple images in an SQL instance.

Before You Begin

Before running the procedures in this chapter, ensure that:

- The database is not in Emergency mode.
- Turn off the SLA options **Run Schedule** and **Expire Data** for the application's policy template.
- Wait for running jobs to finish.

Note: The Restore operation cannot be performed from a remote Actifio Appliance. However, you can restore with a remote-dedup image on the source Actifio Appliance.

Microsoft SQL Server Database Restore Overview

The Restore function initiates all copy data recovery options for an Actifio Appliance and reverts the production data to a specified point in time. Restoring replaces the original production application data with the selected point-in-time image. This restoration results in the loss of all current application data as the application will be restored to its status at the point-in-time when the image was created. This operation cannot be undone.

Note: Actifio provides the flexibility to restore Microsoft SQL Server databases to the original Microsoft SQL Server or to an alternate server. To restore to an alternate server, the Actifio Connector must be installed on the alternate server before initiating the restore operation.

Restore operations are typically performed to restore a database to a valid state after a massive data corruption or storage array failure. The amount of time required to complete a restore operation depends on the amount of data involved.

Databases that use the Microsoft SQL Server Full Recovery Model can use a single policy to capture both the database and its logs. Such a database can be recovered to any point in time by rolling its logs forward. If you restore the database through AGM by specifying Restore With Recovery, the SQL Server database will be restored and brought online without applying logs.

AGM supports the following common use cases when restoring Microsoft SQL Server databases and instances:

- Restore production data: If a production database or instance has become corrupted but it is still on-line, then perform a restore operation.
- Make image data available on the same server or a replacement server: If the database or instance is no longer available on the server, or if you want to put the database on a new server with the same name, then mount a point-in-time image to the server and then perform a restore operation.
- Use of a virtual application (Virtual Application Mount) when you encounter a corrupted SQL Server database: You can use a Virtual Application Mount of the last known good version of a corrupted SQL Server instance or database as a means to allow users and applications to resume work as soon as possible while a new version of the database is rebuilt.
- Discover and capture SQL Server system databases: You can manage the SQL system databases associated with an SQL user database by using a single policy template and resource profile to capture them in a consistency group. This minimizes the consumption of system resources (VDisks) and reduces the number of jobs required to capture data.

Note: Do not use this procedure to restore databases that are members of an AlwaysOn Availability Group. Use the Clone capability to perform parallel clones on all nodes in the SQL AG. See [Chapter 8, Cloning SQL Server Databases](#).

Restoring Microsoft SQL Instances and Databases

Note: If you have copies of multiple SQL Server databases on a single volume, do **not** perform this procedure. It may result in an unintentional data loss as the contents of the entire volume get overwritten during restore of the volume. Instead, Clone the database to another host as detailed in Chapter 8, Cloning SQL Server Databases.

This is the simplest and most common restore scenario. In this case, you restore SQL Server Instances or selected SQL databases from a previous image to the original database server. The database needs to be online for this type of restore. If the database is not online, the restore operation will fail during database validation.

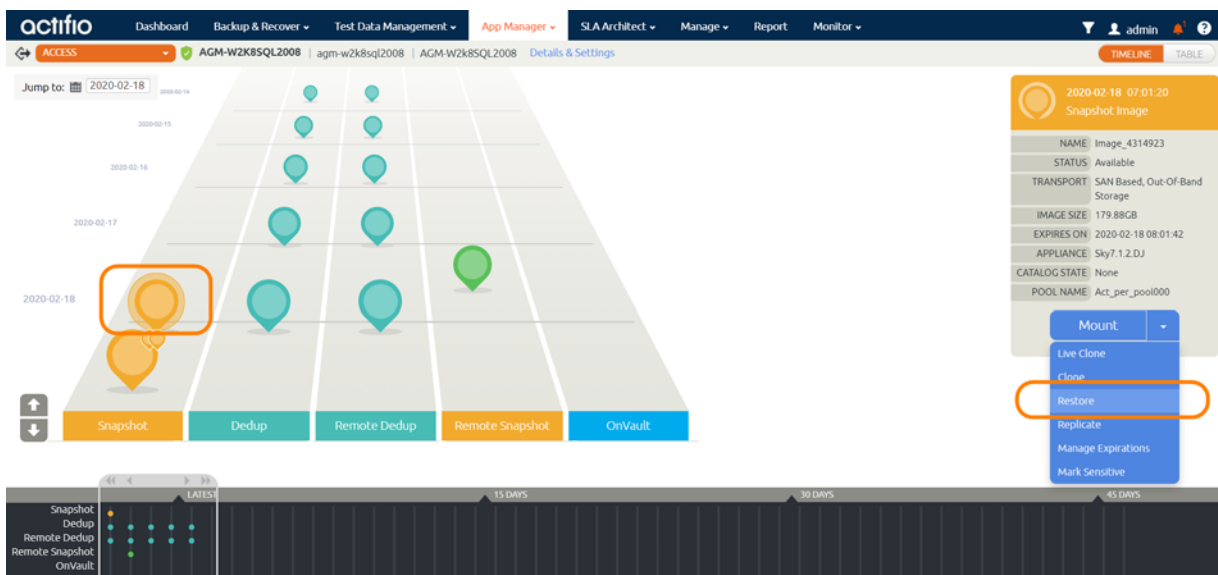
Note: The Restore operation cannot be performed from a remote Actifio Appliance. However, you can use a remote dedup image for a restore operation or a Clone operation on the source Actifio Appliance.

To run this procedure:

- The Microsoft SQL Server database must be online. If the database is not online, the restore operation will fail during database validation.
- The restore operation cannot be performed by AGM from a remote Actifio Appliance. However, you can restore with a remote-dedup image on the source Actifio Appliance.
- Remove SLA management of the SQL Server database, and wait for running jobs to finish.

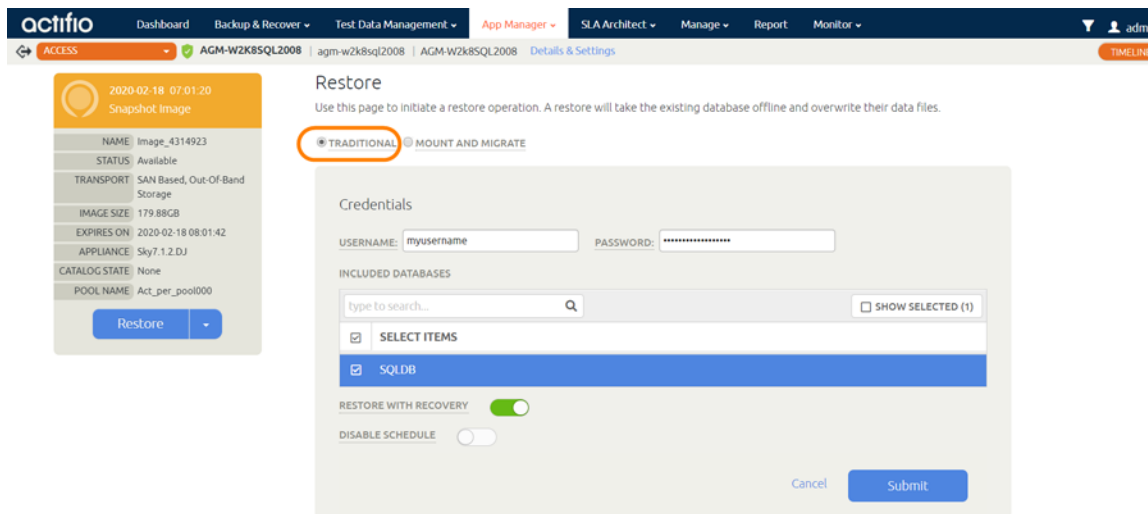
To restore a the SQL Server database(s) or instance:

1. Open the App Manager to the **Applications** page.
2. Right-click the Microsoft SQL Server database that has the image that you plan to restore and select **Manage SLA** from the drop-down list. The Manage SLA page opens.
3. From the Apply dropdown in the upper right corner, disable the SLA.
4. Back at the Applications list, right-click the Microsoft SQL Server database to restore and choose **Access** from the drop-down list. The Access page opens listing captured images in the Timeline ramp view. Image types that support a Restore operation include Snapshot, Dedup, Remote Dedup, and Remote Snapshot (Dedup Async and StreamSnap images).



The background differentiates snapshot images containing an SQL server database with transaction log files, and also illustrates the restore range time period for the logs

5. Select the image, then select **Restore** from the list of operations. The Restore page opens.



6. Select **Traditional** for this restore operation.

Note: In most cases, Mount and Migrate is a better option, as it results in near-zero downtime. Mount and Migrate restores are detailed in [Chapter 7, Mounting and Migrating SQL Data](#).

7. If the selected database does not have logs, the Restore page does not show roll forward options. If the SQL Server database was managed with a Log Protection SLA template, and logs are available with the image, you can:
 - o Specify to roll forward using either User Time or Host Time. You can base the dates and times on User Time or Host time. User Time is relative to the viewer of the current screen. Host time is relative to the system that hosts the data to be restored.
 - o Use the Calendar tool to select a date from which to initiate the restore operation.
 - o Use the Restore Range slider to select a specific point in time to restore the database. Slide the slider tool all the way to the left to restore just the SQL Server database.
8. Select a single volume or multiple volumes to restore. By default all the volumes are selected.
9. Check the Restore With Recovery check box if you do *not* intend to roll the logs. Restore with recovery brings the restored database online. Once online, logs cannot be applied.
10. Click **Submit**. A warning dialog opens. Read it and then enter **DATA LOSS** to confirm. The restore job starts. You can verify that the restore operation is successful by viewing the job status in System Monitor. When the image is restored, AGM creates a new database image populated with data copied from the selected point-in-time image.

Restoring a SQL Server Database to a Different Host

If the original database has been removed, or if the old database server is being replaced with a new server, then use a Clone operation as detailed in [Chapter 8, Cloning SQL Server Databases](#).

Restoring SQL Server Databases in a Consistency Group

Use caution when restoring databases in a consistency group. When you restore databases in a consistency group, all databases in the group are overwritten. If you do not want to overwrite all databases in a consistency group, clone a single database: see [Chapter 8, Cloning SQL Server Databases](#).

For an SQL Server Failover instance, the database is always restored to the active node. VDP mounts the backup image to the active node and performs the restore operation on the node. For SQL Server Availability Groups, the restore is also performed on the active node.

Restoring SQL System Databases

AGM can discover and capture Microsoft SQL system databases just like SQL Server user databases as described in [Restoring Microsoft SQL Instances and Databases](#) on page 55.

As a best practice, capture SQL system databases separate from their associated SQL user databases. This is because SQL system databases are updated less frequently than their associated SQL user databases and can be captured with a much less aggressive schedule.

To manage the SQL system databases associated with an SQL user database, use a single policy template and resource profile to capture them in a consistency group. Doing so minimizes the consumption of system resources (VDisks) and reduces the number of jobs required to capture data.

For example, a single consistency group can capture the SQL system databases:

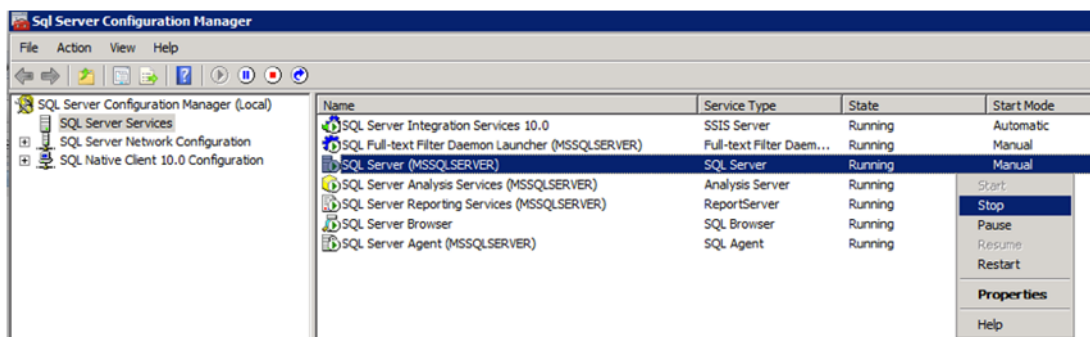
- master
- msdb
- model

To restore an SQL system database, you must first mount the last known good version of the SQL system database consistency group, then use a copy operation to copy the good SQL Server database .mdf and .ldf files to the source SQL server that hosts the corrupt SQL system database.

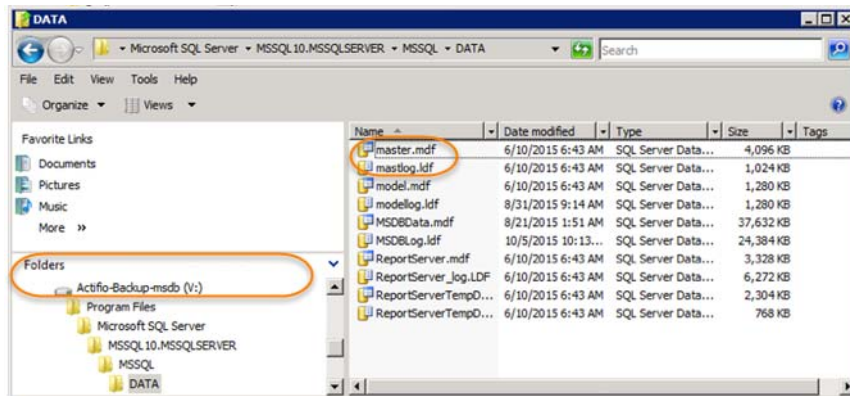
1. From the App Manager Applications list, select and mount the last known good image of the consistency group.




2. From the SQL instance, use either the SQL Server Configuration Manager or the Services MMC to stop the source SQL Server database:



- Using Windows Explorer or some other means, navigate to the mounted SQL system database consistency group:



- Copy the mounted .mdf and .ldf files for the database being restored.
- Using Windows Explorer or some other means, navigate to the source SQL Server database:



- Paste the .mdf and .ldf files into the source SQL Server database.

Use the following sample query to show file locations for databases:

```
SELECT name, physical_name AS current_file_location FROM sys.master_files
```

From the SQL instance, use either the SQL Server Configuration Manager or the Services MMC to Restart the source SQL Server database.

Restoring to an SQL Server Cluster

For an SQL Server Failover instance, the database is always restored to the active node. Actifio VDP mounts the backup image to the active node and performs the restore operation on the node.

For SQL Server Availability Groups, the restore is also performed on the active node. See [Restoring Members of an SQL AlwaysOn Availability Group](#) on page 59.

10 Restoring Members of an SQL AlwaysOn Availability Group

This chapter details how to restore both the primary and secondary members of an AAG to a specific point in time. During the restore process, the primary SQL Server database is overwritten with a backup image and the secondary copies are restored to the same point in time as the restored primary copy.

1. [Identifying the Last Known Good Image of the SQL Server Database](#) on page 59
2. [Restoring the Database on the Primary AAG Node](#) on page 59
3. [Synchronizing Secondary Databases to the Restored Primary Database](#) on page 60
4. [Rebuilding the SQL AlwaysOn Availability Group](#) on page 62

This chapter also provides information on error messages you can encounter during this process. See [Error Messages](#) on page 63 for details.

Before restoring, determine if the database was replicated to another site and if it was, was it replicated using StreamSnap or Dedup-Async. How the database and its logs were replicated will determine how you will recover and restore a database at a secondary site.

You need the logs to bring both the primary and secondary databases up to the same point in time. How the logs were replicated will determine how they will be applied:

- **StreamSnap:** If the database was replicated using StreamSnap and the logs were replicated using the Protect Logs with StreamSnap Technology option, then both the primary and secondary database copies can have the same date specific logs applied.
- **DAR:** If the database was replicated using Dedup-Async Replication (DAR), then the secondary will need a log within a specific date range to get it as close as possible to the primary version.

Identifying the Last Known Good Image of the SQL Server Database

The first step is to learn which of your backup images is the last good image. Performing Virtual Application Mounts of the most recent images and check them. You can mount multiple images simultaneously. Virtual Application Mounts are detailed in [Mounting an SQL Server Database as a New Virtual Database](#) on page 31.

Restoring the Database on the Primary AAG Node

After you have identified the last known good version of the primary database, use AGM to restore the primary production database from that image. Restore operations are detailed in [Chapter 9, Restoring SQL Server Databases](#). When performing the restore, ensure that **Restore with recovery** is ON. The time required to restore the database depends on its size and the capabilities of your environment.

Caution! DO NOT perform log backups on the primary database after restore. This prevents you from applying logs to the secondary databases and prevent the secondary databases from joining the AAG.

Synchronizing Secondary Databases to the Restored Primary Database

For this step, consider the production database on the primary AAG node to be the primary database, and all other databases in the AlwaysOn Availability Group to be secondary databases.

After the primary database has been restored to the original location, it is out of sync with any secondary databases in the AlwaysOn Availability Group. The next step is to get them back in sync.

- Secondary AAG databases that reside locally and are intact can be recovered from the primary database. See [Recovering the Primary From a Non-Corrupt Local Secondary](#) on page 60.
- Secondary databases that are remote may take an unacceptably long time to synchronize. For remote databases, it may be much faster to replace them with an Actifio Clone image:
 - Secondary databases that are captured by an Actifio Production to Mirror policy and reside on a remote Actifio Appliance: see [Restoring a Secondary SQL Server Database From an Actifio Mirror Copy](#) on page 60.
 - Secondary databases that are captured by an Actifio Production to Dedup DR policy and reside on a remote Actifio Appliance: see [Restoring a Secondary SQL Server Database From an Actifio Dedup DR Copy](#) on page 61.

Recovering the Primary From a Non-Corrupt Local Secondary

If both the AAG primary and secondary databases reside locally and if the secondary database is intact, you can recover the primary database from the secondary. To perform a recovery of the primary SQL Server database from a non-corrupt local secondary SQL Server database:

1. After the primary's restore operation has finished, via SQL, remove the database from the AAG:

```
use master
go
alter availability group [AAG-Name] remove database [DATABASENAME];
```
2. Drop the secondary databases:

```
drop database [DATABASENAME];
```
3. Join the replica database to AAG in **Full synchronization**. See [Rebuilding the SQL AlwaysOn Availability Group](#) on page 62 for details.

Restoring a Secondary SQL Server Database From an Actifio Mirror Copy

This approach to restoring a secondary database can be used if the primary database is captured by an Actifio Appliance and then replicated to another Actifio Appliance via a Production to Mirror policy.

You can restore a primary database from a second Actifio Appliance via a Production to Mirror policy if:

- The replicated Mirror copy does not have the same issue(s) as the corrupt primary version.
- You have database log files that were replicated with either:
 - **StreamSnap** along with the database and are the exact logs to use for the restored primary.
 - **Dedup-Async** replication and have a date/time stamp that falls within the range of the log files for the restored primary.

If these conditions can NOT be met, then you must restore the database from a Dedup DR version. See [Restoring a Secondary SQL Server Database From an Actifio Dedup DR Copy](#) on page 61 for details.

To restore a secondary database from an Actifio Mirror copy:

1. Log in to the Actifio Appliance that manages the Mirror copy of the database.
2. Via the App Manager, access the image and view the date time stamp of the Mirror copy.

Note: StreamSnap Mirror policies replicate the production snapshot and its logs. The database and logs for both the production copy and Mirror copy will have the same date/time stamp.

3. For StreamSnap copies, select the version with the same date/time stamp as the production version.
4. Via the Application Manager, Clone the Mirror copy as described in [Chapter 8, Cloning SQL Server Databases](#).
This operation will restore the database in-place and leave it in “restoring” mode (a non-operational state).

Caution: Do not recover the secondary database.

5. When both the primary and secondary databases have been restored and the appropriate logs have been applied, use SQL Studio to recover the primary database. See [Rebuilding the SQL AlwaysOn Availability Group](#) on page 62.

Restoring a Secondary SQL Server Database From an Actifio Dedup DR Copy

This approach to restoring a secondary database can be used if the primary database is captured by an Actifio Appliance and then replicated to another Actifio Appliance via a Dedup DR policy.

Using a Dedup DR copy of the primary database to restore an AAG secondary is best suited for AAGs where:

- An Actifio Mirror copy of the primary database is not available
- Dedup DR versions of the database reside on a second Actifio Appliance

Note: Mounting dedup images requires re-hydration of the image into Snapshot Pool. The size of the image will determine the time required to re-hydrate and the space consumed.

To restore a secondary database from an Actifio Dedup DR copy:

1. Log in to the Actifio Appliance that manages the Dedup DR copy of the database.
2. Select the Dedup DR version of the database that has the same point in time as the last known good version you mounted in [Identifying the Last Known Good Image of the SQL Server Database](#) on page 59.
3. Clone the database image as described in [Chapter 8, Cloning SQL Server Databases](#), and leave it in **restoring mode**.
This operation will restore the database in-place and leave it in “restoring” mode (a non-operational state).


Caution: Do not recover the secondary database.

4. When both the primary and secondary databases have been restored and the appropriate logs applied, use SQL Studio to recover the primary database. See [Rebuilding the SQL AlwaysOn Availability Group](#) on page 62.

Rebuilding the SQL AlwaysOn Availability Group

The restore operation removes databases from an AAG. How the database was restored will determine how you will recover the primary database and rebuild the AAG.

1. From SQL Studio, select the AAG from which the primary database was removed:

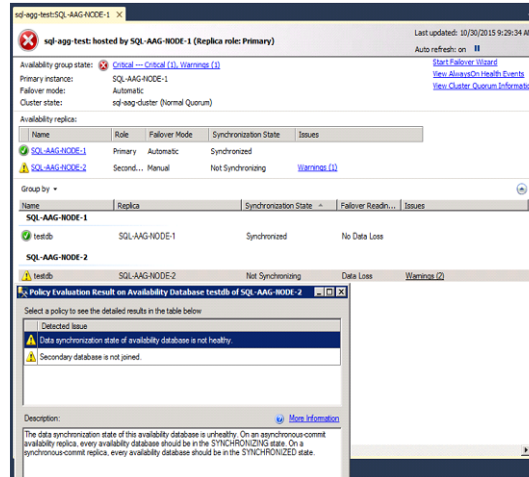


Select AAG


2. When prompted, select the primary database and click **Next**. You will be prompted to select a synchronization method. Select:
 - o **Full**: if the secondary copy of the database was intact and local to the primary database.
 - o **Join only**: if you used method [Restoring a Secondary SQL Server Database From an Actifio Mirror Copy](#) on page 60 or [Restoring a Secondary SQL Server Database From an Actifio Dedup DR Copy](#) on page 61, the Join only operation will form the primary and secondary databases back into an AAG.
3. Monitor the progress of the operation.
4. Once the operation is complete and both the primary and secondary databases are in sync, re-enable the Actifio capture jobs for the AAG.

Error Messages

If you attempt to rebuild an AAG from an Actifio Mirror copy that does not fall within the required range as specified in [Restoring a Secondary SQL Server Database From an Actifio Mirror Copy](#) on page 60, you will encounter one or more of the following errors:

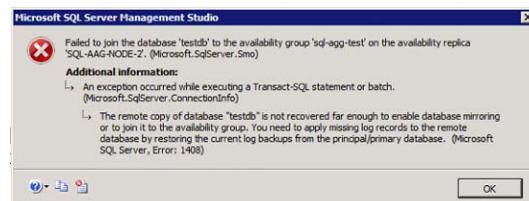



Error Example



Error Example

If logs have been backed up on the primary database after restore, you will not be able to apply the log backup to the secondary and the secondary cannot be joined to the AAG. Then restores and Join AAG operations will fail with:



Error Example

