
An SQL Server DBA's Guide to IBM InfoSphere Copy Data Management

Contents

Chapter 1 - Introduction	1
IBM InfoSphere Data Virtualization	1
Capturing SQL Server Data	2
Capture Mechanisms	3
VMware API Calls	3
The VDP Connector	3
Capturing Microsoft SQL Server Data	4
Capturing SQL Server Database Logs	4
Replicating Logs	4
Resizing a Database Log's Staging VDisk	5
SQL Server Data Capture Options	5
Replicating SQL Server Data	7
Production to Mirror Policy Replication	7
Dedup Backup to Dedup DR Policy Replication	8
Snapshot to OnVault Policy Replication	8
Accessing SQL Server Data	9
Workflows to Automate Access to SQL Server Data	10
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	13
Credentials for Mounting an SQL Server Database as a Virtual Application	14
Chapter 3 - Applying Policy Templates and Resource Profiles	15
Protecting Microsoft SQL Server Instances and Databases	15
Protecting SQL Server Databases in External Storage Pools	17
Protecting SharePoint Data on a Microsoft SQL Server	17
Configuring Application Settings for Microsoft SQL Server Databases	18
Database Log Protection in an SLA Policy	19
Configuring Advanced Settings: Policy Settings Overrides	20

Chapter 4 - Mounting a Microsoft SQL Server Database for Recovery	25
Chapter 5 - Mounting an SQL Server Database as a New Virtual Database	27
Chapter 6 - Cloning SQL Server Databases	31
Chapter 7 - Mounting Encrypted SQL Data	33
Determining if SQL TDE is Enabled	33
Troubleshooting SQL Server Encryption.....	35
SQL Server Master Key, Encryption Certificate, and Password Procedures	36
Chapter 8 - Restoring SQL Server Databases	37
Microsoft SQL Server Database Restore Overview	38
Restoring Microsoft SQL Instances and Databases	39
Restoring a SQL Server Database to a Different Host	40
Restoring SQL Server Databases in a Consistency Group.....	40
Restoring SQL System Databases.....	41
Restoring to an SQL Server Cluster.....	42
Chapter 9 - Restoring Members of an SQL AlwaysOn Availability Group	43
Identifying the Last Known Good Image of the SQL Server Database	43
Restoring the Database on the Primary AAG Node	43
Synchronizing Secondary Databases to the Restored Primary Database	44
Recovering the Primary From a Non-Corrupt Local Secondary	44
Restoring a Secondary SQL Server Database From an IBM InfoSphere Mirror Copy.....	44
Restoring a Secondary SQL Server Database From an IBM InfoSphere Dedup DR Copy	45
Rebuilding the SQL AlwaysOn Availability Group	46
Error Messages.....	47
Chapter 10 - Using Mounts to Create SQL AlwaysOn Availability Groups	49
Creating an SQL Server AAG in an IBM InfoSphere Snapshot Pool.....	49
Creating an SQL Server AAG Outside of An IBM InfoSphere Snapshot Pool.....	49
Creating the New SQL Server AlwaysOn Availability Group.....	50

1 Introduction

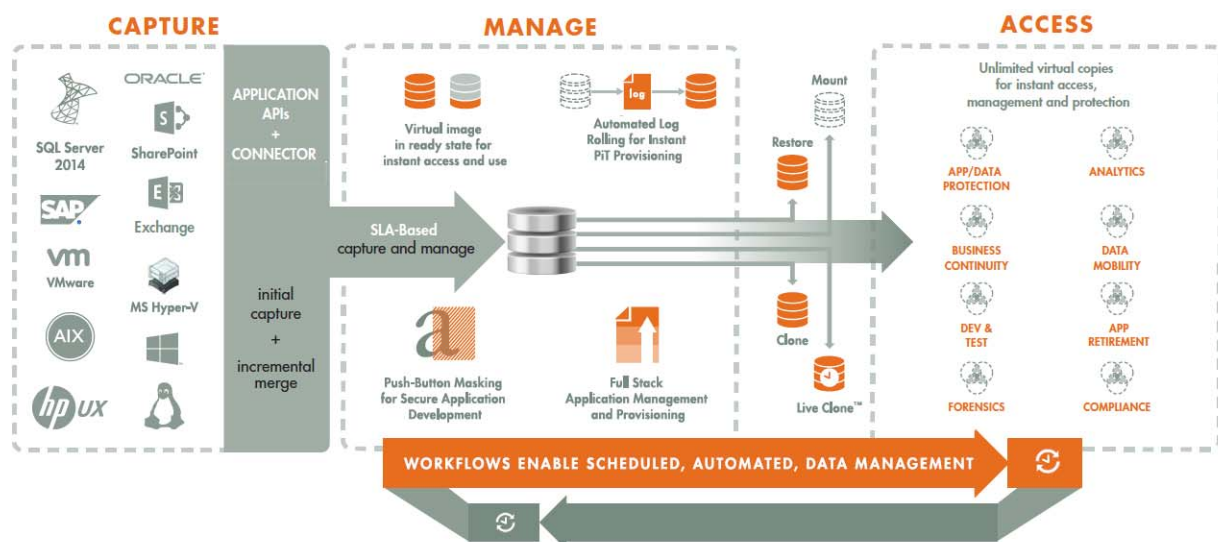
This chapter provides a high-level overview of basic IBM InfoSphere concepts and procedures used to capture and access Microsoft SQL Server databases:

- [IBM InfoSphere Data Virtualization](#) on page 1
- [Capturing SQL Server Data](#) on page 2
- [Capture Mechanisms](#) on page 3
- [Capturing Microsoft SQL Server Data](#) on page 4
- [Capturing SQL Server Database Logs](#) on page 4
- [Replicating Logs](#) on page 4
- [Resizing a Database Log's Staging VDisk](#) on page 5
- [SQL Server Data Capture Options](#) on page 5
- [Replicating SQL Server Data](#) on page 7
- [Accessing SQL Server Data](#) on page 9
- [Workflows to Automate Access to SQL Server Data](#) on page 10

IBM InfoSphere Data Virtualization

An InfoSphere VDP 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.

IBM InfoSphere 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.



Capture, Manage, and Access Application Data

Application data is captured at the block level, in application native format, according to a specified 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.

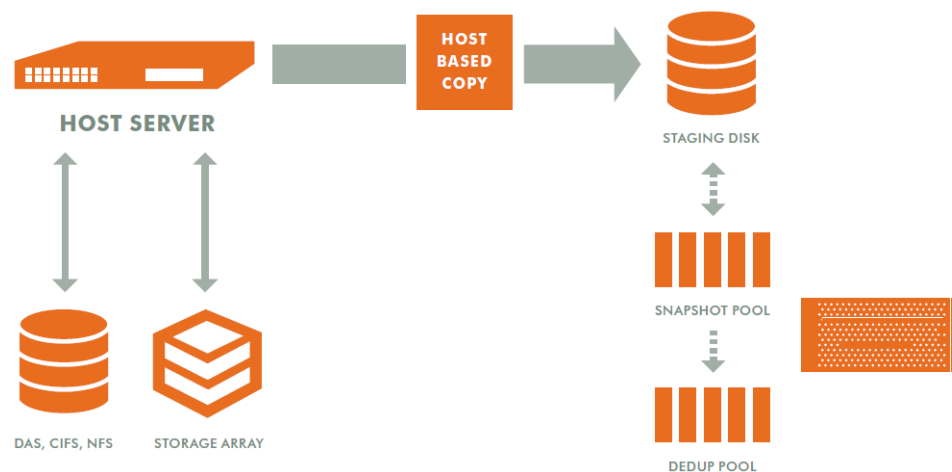
Capturing SQL Server Data

The IBM InfoSphere user interfaces allow you to capture SQL:

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

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

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



IBM InfoSphere Data Capture

When capturing data:

- A staging disk is automatically created and mounted on a server.
- An initial full copy is made to the staging disk. Subsequent copies consist only of changed blocks.
- The staging disk is unmounted from the server.
- A snapshot of the staging disk is made on the InfoSphere VDP Appliance.

IBM InfoSphere offers an alternate configuration where production data storage is controlled by an InfoSphere VDP 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

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

IBM InfoSphere has two methods of capturing Windows data with VSS. The first is to leverage VMware-level captures, which is integrated with the VMware Tools VSS provider for consistent data captures. The second is IBM InfoSphere proprietary and leverages the built-in Microsoft VSS provider. This comes with all the safeguards for both performance and space management that Microsoft has built, and is in-line with Microsoft's best practices. IBM InfoSphere'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.

Customers with the most highly transactional databases are able to take IBM InfoSphere snapshots of their databases without an observed performance impact to their users. IBM InfoSphere 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 IBM InfoSphere to snapshot their Exchange DAG primary nodes without triggering failovers.

VMware API Calls

An InfoSphere VDP Appliance can take advantage of VMware API for data protection (VADP) calls to capture an entire virtual server. Specifically, the API calls can:

Perform change block tracking: Makes an initial full snapshot of the entire VM, then going forward only snapshots the changes to the database thereby enabling IBM InfoSphere's "incremental forever" capture strategy.

Quiesce applications: Ensures application consistency during capture.

When an entire virtual server is managed, a fully functional virtual server (operating system, applications, and their data) is captured. Having a copy of the entire virtual server guarantees that the data can be accessed fast and without issues. Because the image presented is a fully functional virtual server, if needed, it can be started and run from an InfoSphere VDP Appliance directly and then optionally migrated to a new, permanent location. Virtual machines and their applications can be grouped and captured with a single SLA Template Policy.

The VDP Connector

The VDP Connector is used to capture with granularity at the Microsoft SQL Server database level. The VDP Connector is a small-footprint, lightweight service that can be installed on either virtual or physical servers. The VDP 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 VDP Connector allows you to capture 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 VDP 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 thereby ensuring a consistent point in time capture of data.
- Identifies changes to database data for IBM InfoSphere'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 if they are 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. For Microsoft SQL Server failover instances, discovery should be performed on all cluster nodes to ensure both clustered and non-clustered instances and filesystems are discovered.
3. Define IBM InfoSphere 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 IBM InfoSphere 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 Advanced Options. 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 IBM InfoSphere 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 InfoSphere VDP Appliance. You can use the logs at the remote site for any database image within the retention range of the replicated logs. Log replication uses StreamSnap technology to perform the replication between the local and remote InfoSphere VDP Appliances; the replication goes directly from the snapshot pool of the local InfoSphere VDP Appliance to the remote snapshot pool.

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 InfoSphere 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 InfoSphere 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.

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

Note: Log replication does not occur until an SQL Server database has been protected and the database is replicated to the remote InfoSphere VDP Appliance.

For details on creating policies and using policy options, see **Planning and Developing Service Level Agreements**.

Resizing a Database Log's Staging VDisk

The physical space required to accommodate a database's logs is automatically managed by the InfoSphere VDP Appliance. At a minimum, the InfoSphere VDP Appliance will evaluate 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. This setting is from 5 to 100 percent (default is 50%).
- **Estimated Change Rate** - Defines the daily change (in percent), which allows the InfoSphere VDP Appliance to better calculate the size of the staging disk needed to hold logs. This setting is from 0 to 100 (default is 10%).
- **Compress Database Log Backup** - Instructs the source database to compress its logs before capture by the InfoSphere VDP Appliance. The database server performs log compression during log backup (default is Enabled).

For details on database log advanced settings, see the [online help](#).

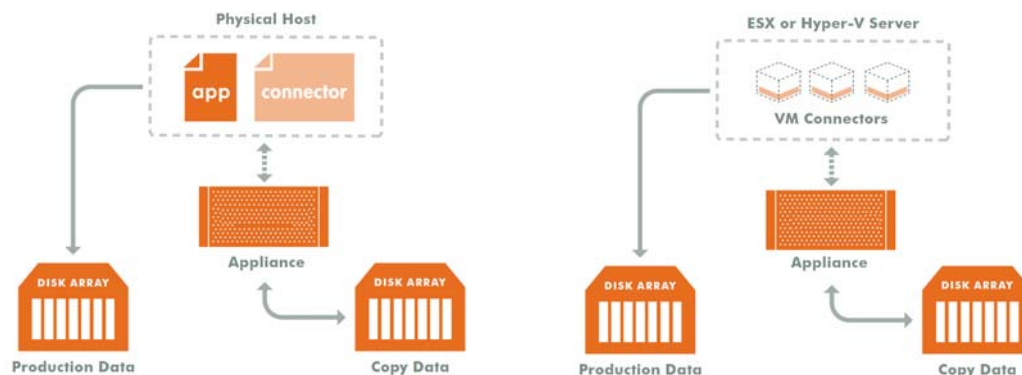
SQL Server Data Capture Options

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 6
- [Capturing a VM's Databases and Boot Volume](#) on page 6
- [Capturing Entire VMware and Hyper-V VMs](#) on page 6

Capturing Instances, Individual Databases, and Groups of Databases

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



Capturing Individual or Groups of Databases

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 IBM InfoSphere capture job. Databases in an instance are quiesced and captured together with a single IBM InfoSphere SLA.

If IBM InfoSphere's database and log capture technology is 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 IBM InfoSphere user interface with a single action.

Individual members of an instance can be accessed by a 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 IBM InfoSphere SLA Policy Template and Resource Profile. Membership to a consistency group is done 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.

As the name implies, consistency groups ensure consistent point-in-time capture and recovery across multiple databases. If IBM InfoSphere'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 IBM InfoSphere 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:

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

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

Databases in a clustered failover instance must be discovered from the active node. Once protected, the InfoSphere VDP Appliance 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.

Capturing Entire VMware and Hyper-V VMs

To capture entire VMware VMs, the InfoSphere VDP Appliance takes advantage of VMware APIs. To capture entire Hyper-V VMs, the InfoSphere VDP Appliance uses an VDP Connector installed on the Hyper-V server.



Capturing Entire VMs

When an entire virtual server is captured, the image is of a fully functional server (operating system, applications and their data). This guarantees that the data can be accessed fast and without issues. Because the image presented is a fully functional virtual server, it can be migrated to a new, permanent location.

Capturing whole virtual servers allows groups of virtual servers and their applications to be protected with a single SLA Policy Template.

Replicating SQL Server Data

Data can be replicated to a second InfoSphere VDP 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. IBM InfoSphere 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 InfoSphere VDP 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 InfoSphere VDP Appliances is performed using 1024-bit certificates.

Replication is controlled by IBM InfoSphere Policy Template policies:

- Production to Mirror policies have several options to replicate data to a second InfoSphere VDP Appliance. For details see [Production to Mirror Policy Replication](#) on page 7.
- Dedup Backup to Dedup DR policies use a fixed, IBM InfoSphere proprietary replication engine to replicate data to a second InfoSphere VDP Appliance. In addition Dedup Backup to Dedup DR policies allow you to replicate data to two locations. For details see [Dedup Backup to Dedup DR Policy Replication](#) on page 8.
- Production to OnVault policies use a fixed, IBM InfoSphere proprietary engine to transfer data to object storage. For details see [Snapshot to OnVault Policy Replication](#) on page 8.

Production to Mirror Policy Replication

Production to Mirror policies can use one of the following options to replicate data to a second InfoSphere VDP Appliance:

- StreamSnap
- Dedup Async (DAR)

When data is replicated via a Production to Mirror policy, it lands in a Snapshot Pool managed by another InfoSphere VDP Appliance. The Snapshot Pool on which that replicated data lands must be sized accordingly.

StreamSnap

StreamSnap facilitates high-availability by allowing you to keep a remote copy of an SQL database, along with its logs up-to-date and ready for a failover scenario. When a StreamSnap-managed application fails, you mount a failover image of the application from the remote site. When the problem has been resolved, then you can restore the syncback image to the local site with the latest changes and then failback the application to the production site.

StreamSnap replicates snapshots of databases and their logs to a remote InfoSphere VDP Appliance over a high quality bandwidth IP network, which can provide RPOs as low as one hour.

- For VMware VMs, snapshot replication is streamed to the second InfoSphere VDP Appliance in parallel. Streaming of a VMware VM is performed to avoid waiting until the local snapshot job completes before initiating replication.
- For non-VMware VM applications, snapshot replication occurs after the local snapshot job is completed.

Note: StreamSnap replication and local snapshots are integrated to avoid the creation of double snapshots. The InfoSphere VDP Appliance allows you to maintain multiple local snapshots and store local images in the Dedup pool for long-term retention.

Production to Mirror policies that use StreamSnap replication are tied to a specific Production to Snapshot policy. They use the schedule and frequency settings of their associated Production to Snapshot policy.

StreamSnap replication requires a reliable network connection to replicate snapshots to the remote InfoSphere VDP Appliance. The bandwidth required is directly related to the application size (initial copy) and change rate (incremental updates).

Dedup Async (DAR)

Dedup-Async Replication (DAR) is an IBM InfoSphere-proprietary form of replication where initially a full copy of data is replicated to another InfoSphere VDP Appliance, then going forward, the copy is updated with incremental changes. This reduces the amount of data that must be sent over the network and ensures that an up-to-date copy of production data is always present and ready for a recovery operation.

Note: *Production to Mirror policies that use DAR make snapshots of their own. They do not use a snapshot created by another Production to Snapshot policy and they cannot replicate database logs.*

Because the data is deduplicated before it is replicated, less network bandwidth is required. DAR is used to:

- Achieve typical Recovery Point Objectives (RPOs) of 24 hours with 12 and 8 hour RPOs possible.
- Replicate data that is can be efficiently deduplicated.

IBM InfoSphere's Dedup-Async replication:

- Uses existing IP network to replicate data.
- Replicates repeatedly at intervals determined by policy.
- Makes disk management transparent.
- Replicates VMware VMs to a datastore (optional).
- Makes fail-over to a host on the remote site simple.
- Makes synback to the local InfoSphere VDP Appliance simple and non-interruptive to the application's input/output activities.
- Cleans up all remote images created after the application fails back to the local InfoSphere VDP Appliance.

Dedup Backup to Dedup DR Policy Replication

Dedup Backup to Dedup DR replication provides for efficient long-term storage of deduplicated captured data. Dedup Backup to Dedup DR replication is scheduled IP-based replication. Dedup DR is designed for long-term storage of deduped data on a remote InfoSphere VDP Appliance. When data is replicated via a Dedup Backup to Dedup DR policy, it lands in a Dedup Pool managed by another InfoSphere VDP Appliance. The Dedup Pool on which that replicated data lands must be sized accordingly.

The Dedup Backup to Dedup DR replication process begins after the deduplication process completes. A proprietary deduplication-aware replication protocol enables the transmission of only the globally unique blocks.

This approach minimizes the bandwidth required to move data between InfoSphere VDP Appliances.

Finally, Dedup Backup to Dedup DR replication provides the added benefit of allowing data to be replicated to a remote site, then from that remote site to a second remote site. This feature is referred to as multi-hop.

Snapshot to OnVault Policy Replication

The Snapshot to OnVault policy allows you to transfer data to remote storage. The transfer of data is scheduled within the policy and the most recent snapshot taken by the Policy Template will be transferred. An HTTPS connection is used to ensure data security and to minimize network traffic.

Accessing SQL Server Data

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

For Microsoft SQL Server databases that use the Simple Recovery Model, the InfoSphere VDP Appliance 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

IBM InfoSphere administrators can control which users have access to data, IBM InfoSphere features, processes, and resources. In addition, captured data can be defined as sensitive or non-sensitive. IBM InfoSphere users can be granted permission to access sensitive data.

Mounts

The IBM InfoSphere mount function provides instant access to data without moving data. Captured copies of databases can be rolled forward via the IBM InfoSphere user interface and mounted on any database server. An InfoSphere VDP Appliance provides two ways to mount an Microsoft SQL Server database:

- **The standard 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.
- **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 address the unique challenges associated with creating and managing copies of production databases for non-production use. Virtual application mounts are created from the InfoSphere VDP Appliance and do not require manual intervention by database, server, or storage administrators. Virtual application mounts can be used for such things as database reporting, analytics, integrity testing, and test and development.

Note: If you mount an SQL database as a virtual application to an SQL instance, the virtual application cannot be captured along with the other databases in the instance. Virtual databases must be captured separately.

LiveClones

A LiveClone is an independent copy of Microsoft SQL Server data that can be refreshed when source data changes. LiveClones are independent copies of data that can be incrementally refreshed and masked before being made available to users. This allows teams such as development and test to ensure they are working on the latest set of data without having to manually manage the data and not access or interfere with the production environment.

Restores

The restore function 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. IBM InfoSphere's functionality allows you to, from the IBM InfoSphere user interface, restore the database in Restoring Mode and then roll the logs forward to a specific point in time.

If you restore the database through the IBM InfoSphere user interface without specifying Restore with no Recovery, the database will be restored and brought on line without applying logs.

Workflows to Automate Access to SQL Server Data

While SLAs govern the automated capture of production Microsoft SQL Server data. Workflows automate access to the captured Microsoft SQL Server data. Workflows are built with 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 on 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 Microsoft SQL Server data in a LiveClone prior to making it available to users.

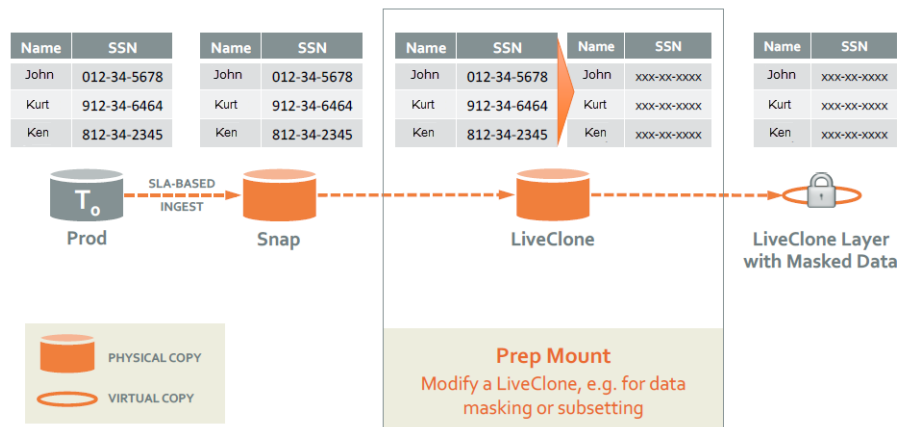
Combining IBM InfoSphere's automated Microsoft SQL Server data capture, access control, with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now, instead of having to wait weeks for DBAs to update test and development environments, users can provision their own environments almost instantly.

For example, an IBM InfoSphere administrator can create an SLA Template Policy that captures Microsoft SQL Server data according to a specified schedule. Optionally, 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 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, via the IBM InfoSphere user interface, 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***.

2 Required SQL Server Roles for the Windows User

Microsoft SQL Server requires specific user roles to perform specific operations. To perform IBM InfoSphere capture, restore, unmount, delete, and Application Aware mount operations on an SQL Server database, you must provide the InfoSphere VDP Appliance 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 Application Aware mount operations from an InfoSphere VDP 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 IBM InfoSphere 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 13
- [Credentials for Mounting an SQL Server Database as a Virtual Application](#) on page 14

Windows Local Admin User

To perform capture, restore, unmount delete, and Application Aware Mounts you must enter 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 IBM InfoSphere capture, restore and Application Aware 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 IBM InfoSphere 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 IBM InfoSphere 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 application's SLA Application Details & Settings.

The screenshot displays the 'Application Details & Settings' window in the IBM InfoSphere Application Manager. The window is titled 'Application Details & Settings' and has a 'Settings Help' link. It contains a table of application details for 'AAGDATA1' and a 'Settings' section. The 'Settings' section has two input fields: 'USERNAME' and 'PASSWORD', which are highlighted with an orange rectangle. The background shows the 'APPLICATION MANAGER' interface with tabs for Applications, Consistency Groups, Logical Groups, Active Images, and Workflows. The 'MANAGE SLA' tab is active, showing a list of SLA templates with 'DAR/Policy' selected.

Application Name: AAGDATA1	
APPLICATION TYPE	SQLServer
HOST	AAG2K19N1N2N3.FERRARI.COM
HOST IP ADDRESS	172.16.22.19
PATH	AAG2K19N1N2N3.FERRARI.COM
OPERATING SYSTEM	Win32
APPLIANCE	Sky901
APPLIANCE IP ADDRESS	172.16.22.22
CLUSTER NODE LIST	2019AAG-Node1,2019AAG-Node2,2019AAG-Node3

Settings

USERNAME

PASSWORD

Cancel Save Changes

Credentials for Capturing Data

Note: Credentials are not required if only databases are being captured.

Credentials for Restoring a Microsoft SQL Server Database

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

APPLICATION MANAGER Filter by Appliances Organizations America/New_York admin

Applications Consistency Groups Logical Groups Active Images Workflows

ACCESS ActifioDB06 sql2012SQA.sqa.actifio.com 172.16.200.100 Details & Settings

2019-05-22 05:08:05 Snapshot Image

NAME: Image_0005519
STATUS: Available
TRANSPORT: SAN Based, Out-Of-Band Storage
IMAGE SIZE: 71.59GB
EXPIRES ON: 2019-05-24 05:09:55
APPLIANCE: Amazon
RECOVERY RANGE: 05-22 05:08 To 05-22 13:00
CATALOG STATE: None

Restore

Restore

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

ROLL FORWARD TIME 2019-05-22 10:00:59 HOST TIME USER TIME

Credentials

USERNAME: PASSWORD:

Select volumes to restore Deselect All Volumes Add All Volumes

C:\

CAPACITY	72 GB
UNIQUE ID	dasvol.C\
VOLUME TYPE	Non-Boot
TARGET	vdisk.fc:5CE511804900

RESTORE WITH RECOVERY ☒

DISABLE SCHEDULE ☐

Cancel Submit

Credentials for Restore Operation

Credentials for Mounting an SQL Server Database as a Virtual Application

An application aware mount mounts an SQL Server database as a virtual application. When performing an Application Aware mount of an SQL Server database from an InfoSphere VDP Appliance, the user must be assigned a role that allows both the ability to mount and unmount (detach) the SQL Server database.

When performing an Application Aware mount, in the Mount dialog box Advanced Options, enter the credentials of a Windows user assigned the proper role(s):

The screenshot shows the 'Mount' dialog box in the IBM InfoSphere VDP Appliance interface. The dialog is titled 'Mount' and has a 'TARGET' dropdown set to 'sql2012SQA.sqa.actifio.com'. The 'LABEL' field is empty. Under 'Application Options', the 'CREATE NEW VIRTUAL APPLICATION' toggle is turned on. The 'ROLL FORWARD TIME' is set to '2019-05-22' at '10:00:59'. The 'SQL SERVER INSTANCE NAME' is 'SQL2012SQA/SQL2012INS' and the 'SQL SERVER DATABASE NAME' is 'SQL2012SQA/SQL2012INSTA92'. The 'MANAGE NEW APPLICATION' toggle is turned off. The 'Advanced Options' section is expanded and highlighted with an orange circle. It contains two toggles: 'RECOVER DATABASE AFTER RESTORE' (turned on) and 'RECOVER USER LOGINS' (turned off). Below these are two text input fields: 'USER NAME' and 'PASSWORD', which are highlighted with an orange rectangle. The left sidebar shows a list of applications, with '2019-05-22 05:08:05 Snapshot Image' selected. The bottom of the sidebar has a 'Mount' button.

Credentials for Mount Operation

3 Applying Policy Templates and Resource Profiles

To protect a Microsoft SQL Server instance or database, you must create SLA Policy Templates and Resource Profiles, and then apply them to databases. Step-by-step best practices for creating SLA Policy templates and resource profiles can be found in the IVGM online help.

A database is first 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. This chapter includes:

[Protecting Microsoft SQL Server Instances and Databases](#) on page 15

[Protecting SQL Server Databases in External Storage Pools](#) on page 17

[Protecting SharePoint Data on a Microsoft SQL Server](#) on page 17

[Configuring Application Settings for Microsoft SQL Server Databases](#) on page 18

[Database Log Protection in an SLA Policy](#) on page 19

[Configuring Advanced Settings: Policy Settings Overrides](#) on page 20

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**.

Protecting Microsoft SQL Server Instances and Databases

IVGM allows you to capture:

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

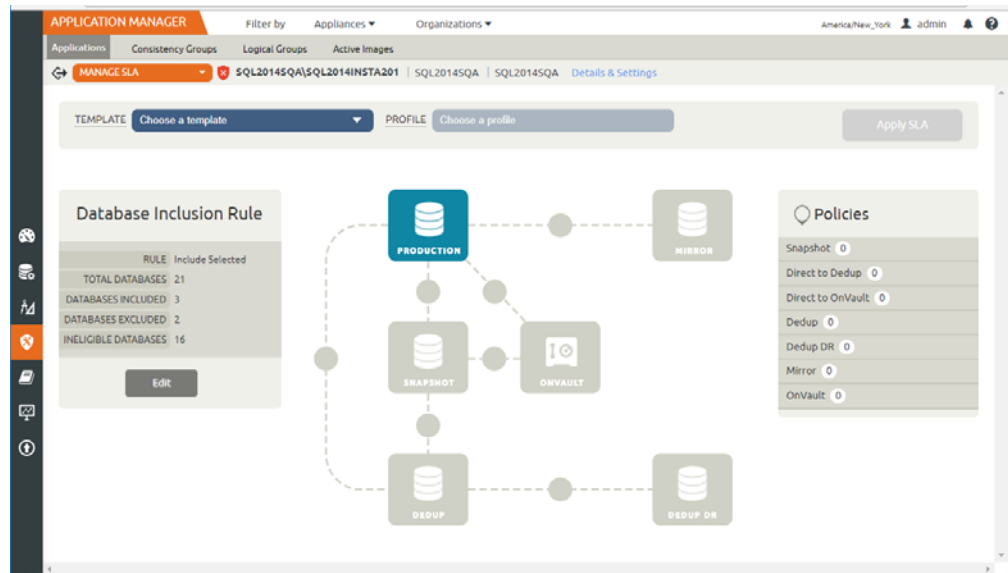
For background information on managing and capturing Microsoft SQL Server databases, see *Managing Microsoft SQL Databases* in the IVGM online help.

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. For details, see [Protecting SharePoint Data on a Microsoft SQL Server](#) on page 17.

To capture a Microsoft SQL Server database:

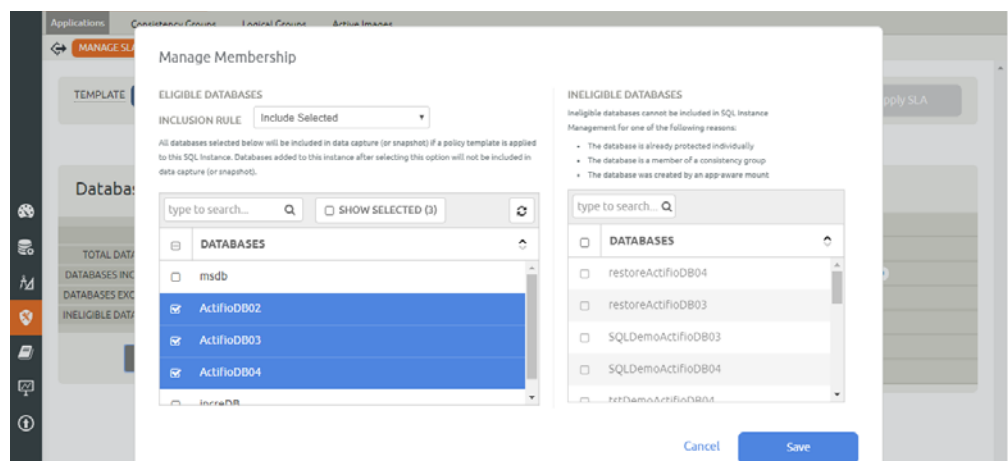
1. From the IVGM left-hand navigation, click the Application Manager icon. The Applications page opens.
2. Select the Microsoft SQL Server database, instance, AAG, or consistency group that you want to capture.

3. Select **Manage SLA** from the drop-down list at the bottom right corner, and the Manage SLA page opens.



The Manage SLA Page

4. 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.
5. 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 18 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 20 for details.
6. To select databases, under Database Inclusion Rule, click **Edit**. The Manage Membership dialog box opens:



Managing Membership

7. From the Manage Membership dialog box, select the databases to capture.
8. Click **Save** and the Manage Membership dialog box closes.
9. 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.

Protecting SQL Server Databases in External Storage Pools

If 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.

Protecting SharePoint Data on a Microsoft SQL Server

SharePoint data is protected by capturing its SQL Server database. When capturing a SharePoint SQL Server database, application advanced settings specific to SharePoint are available.

To configure application advanced settings for a SharePoint application residing on a SQL server:

1. Open the Application Manager to the **Protect** tab.
2. Select **APP**. Under **Local > SharePoint**, select the SharePoint server from the navigation pane.
3. Click the blue **Advanced Settings** link in the upper right corner. The SQL Server Application Advanced Settings page appears.
4. **Do Not Unmap**: Select an unmapping option. Temporary staging disks are mapped to the host and used during data movement. Select whether you want them to remain mapped to the host. By default, LUNs are mapped during the first job and all the subsequent jobs reuse the same mapped LUN.
5. Select whether you want to **Truncate Logs** after every backup. When this option is selected, application related logs are truncated to the current or most recent backup.
6. **Service Access Point IP Address**: Service Access Point is relevant only for SQL server availability groups. Enter a value here to back up from a SQL availability cluster. Specify the IP address of the cluster 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 cluster.
7. Leave **Connector Options** blank unless you are working with IBM InfoSphere Support.
8. Click **Save** to update your changes.

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 the check box to the left of the selection you want to reset, or click **Select options that will revert back to default** to reset all application selections to their default state.

Note: The following list details all of the SQL application details and settings. The actual list displayed will vary by whether an SQL Instance, SQL Cluster, or AAG is selected.

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 would be saved in this directory. Ensure that there is enough free space available on the volume hosting this directory to hold a full database backup.
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.
Connector Options	Leave Connector Options blank unless directed to enter a value by IBM InfoSphere 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 have the option of also capturing 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 advanced settings of the database's snapshot policy. The capture of logs is done without regard to day boundaries, window, or frequency at which 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 also defined in the advanced settings for an SLA Policy. The capture of logs is done without regard to day boundaries, window, or frequency at which its associated database is captured.

Policy Settings

The physical space required to accommodate a database's logs is automatically managed by IVGM. At a minimum, IVGM 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 the VDP appliance. The database server performs log compression during log backup.

You can replicate database logs to a remote VDP 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 VDP 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.

Policy Settings Overrides Valid for SQL Server Instances, Databases, and Consistency Groups

Advanced Setting	Description
Enable File Catalog	When enabled, the VDP Catalog will scan and index the captured data.
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 VDP Catalog.
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. <hr/> <p>Note: For applications managed using the VDP Connector (such as SQL or Exchange databases) 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> <hr/> <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	<p>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.</p> <p>Options are:</p> <ul style="list-style-type: none">Do not truncate/purge log after backupTruncate/purge log after backup

Policy Settings Overrides Valid for SQL Server Instances, Databases, and Consistency Groups

Advanced Setting	Description
Skip Offline Applications in the Consistency Group (For Consistency Group management only)	<p>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:</p> <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	<p>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.</p> <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	<p>Captures the SQL Server database login credentials. When the database is mounted as a virtual application (application aware mount) the virtual database will have all of the login credentials used by the source. Options are Yes or No.</p>
Enable Database Log Backup	<p>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> <hr/> <p>Note: For details on Log Protection, see Database Log Protection in an SLA Policy on page 19.</p> <hr/>
RPO	<p>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.</p>
Log Backup Retention Period (In Days)	<p>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.</p>

Policy Settings Overrides Valid for SQL Server Instances, Databases, and Consistency Groups

Advanced Setting	Description
Replicate Logs (Uses StreamSnap Technology)	<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: Log replication does not occur until an Oracle or SQL Server database has been protected and the database replicated to the remote VDP appliance.</p> <hr/>
Log Staging Disk Growth Size (In Percent)	<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>
Estimated Change Rate	<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>
Compress Database Log Backup	<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 IVGM. 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>
Enforced Retention	<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: Enforced Retention cannot be overridden on a per-application basis. The option does not appear on the "Policy Overrides" page.</p> <hr/>

Policy Settings Overrides Valid for SQL Server Instances, Databases, and Consistency Groups

Advanced 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 UDSAgent: Specified start path does not exist. If no start paths are specified, this option has no effect. Options are Yes or No.</p> <hr/> <p>Note: 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.</p> <hr/>
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: By default, the script timeout values are 1. If a script timeout is not specified, the value will be blank.</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>

4 Mounting a Microsoft SQL Server Database for Recovery

This chapter describes how to mount captured Microsoft SQL Server database data. You can mount Microsoft SQL:

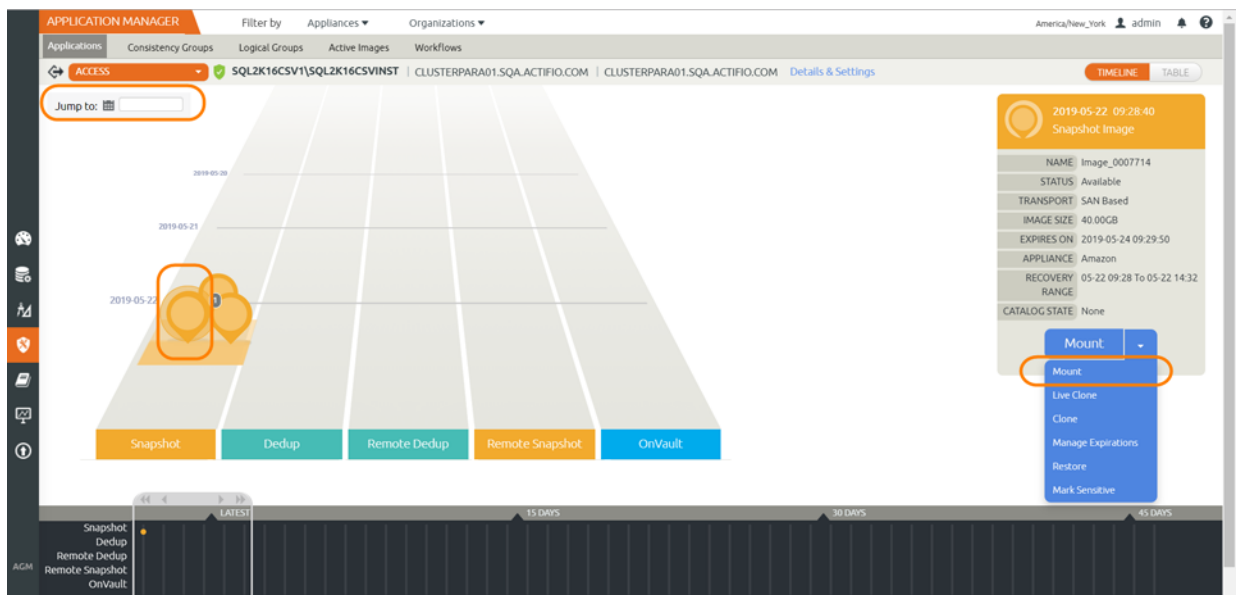
- Instances
- System databases
- User databases
- Consistency groups
- Individual members of a consistency group
- Primary database of an AlwaysOn Availability Group

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

Options presented in the wizard will 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.

With this option 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 IVGM to the **Application Manager**.
2. Right-click an SQL instance, user database, system database, cluster, or availability group and select **Access**. The filters in the left-hand pane can make it easier to find the database that you need.
3. On the runway of images, select the specific image to be mounted. On the right side, select **Mount**.

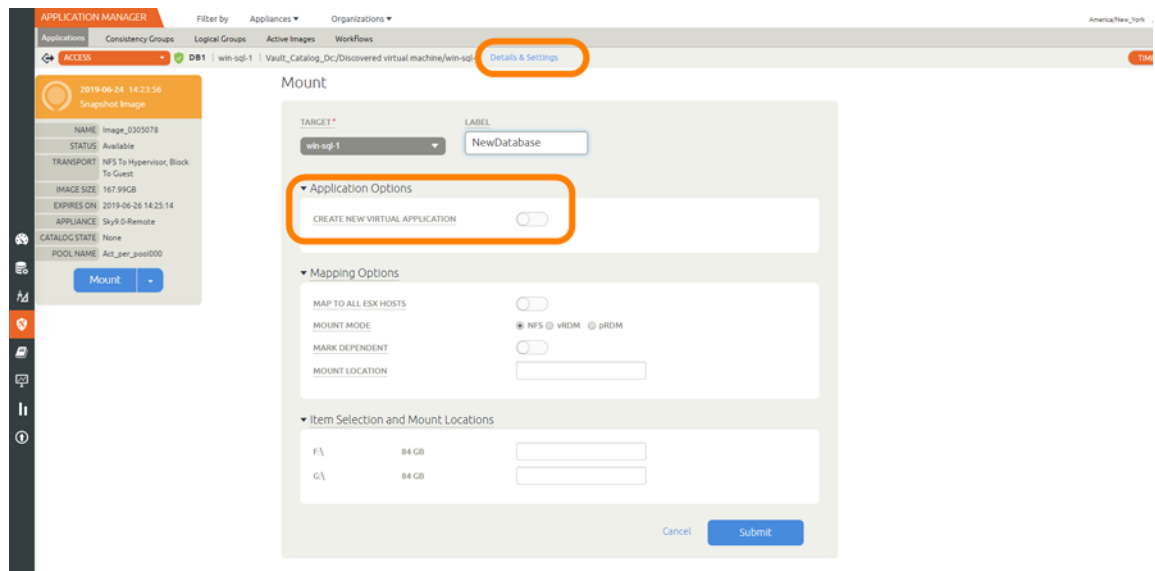


Selecting a SQL Image to Mount

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

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 InfoSphere VDP Appliances. Once the entire image is copied, then the mount is performed from the snapshot pool.



Initial Mount Screen for an SQL Server Database

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 Database**. (To mount the Microsoft SQL data as a virtual database, see [Chapter 5, Mounting an SQL Server Database as a New Virtual Database](#).)
7. Fill in **Mapping Options** as needed for this new database.
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.

5 Mounting an SQL Server Database as a New Virtual Database

An Application Aware 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. An Application Aware 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 InfoSphere VDP Appliance.

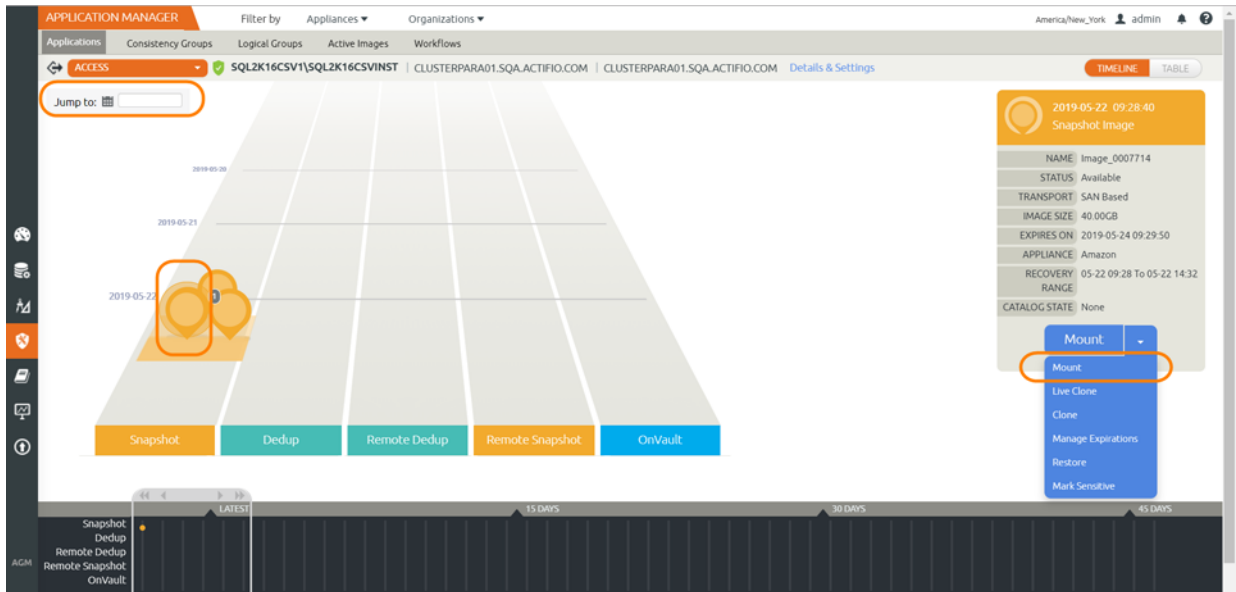
Note: A database mounted to a protected SQL Instance as a virtual application will not be protected with the SQL Instance.

If you want to mount just the Microsoft SQL data for database recovery, see [Chapter 4, Mounting a Microsoft SQL Server Database for Recovery](#).

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

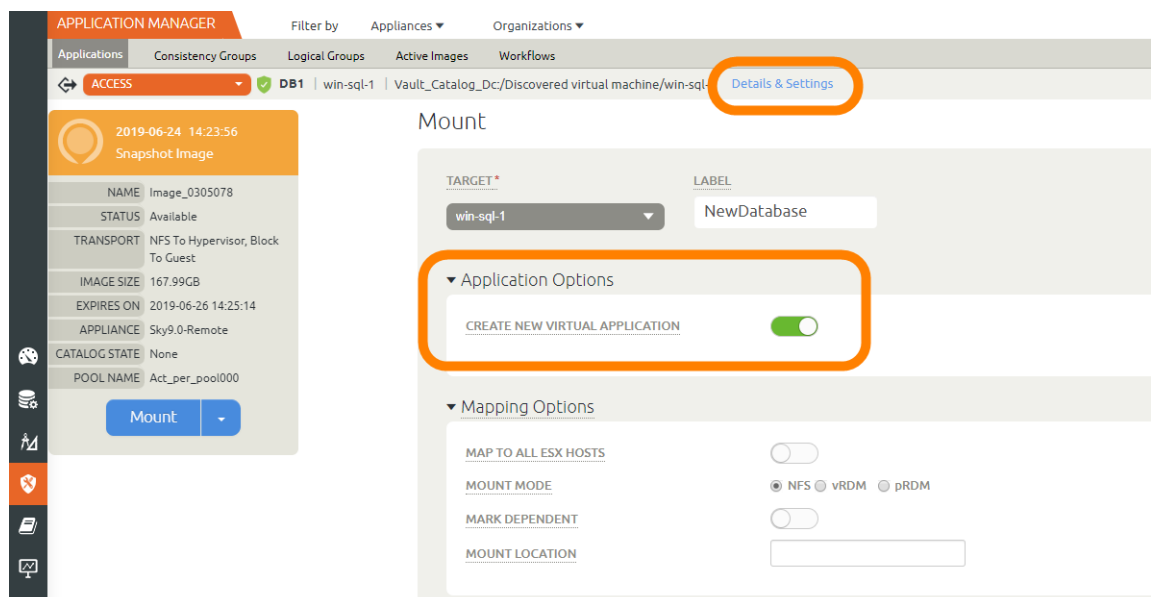


Note: You can use the calendar widget in the upper left corner to narrow the number of backup images available.

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

Note: OnVault images can be mounted very efficiently from InfoSphere VDP 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 captured database was captured along with its logs, the App Options dialog box will provide an option to rolls 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 Application Aware mount will be a new, virtual database. If you want to protect the new virtual database, then select **Manage New Application** and select the template and profile to use.

The application aware 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 forget and 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 InfoSphere VDP 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. Check **Recover Database After Restore**. Doing so leaves the database in a state where if logs are available they can be applied to bring the database to a specific point in time.
13. Check **Recover User Logins** to back up the master database. This records all the system-level information for a SQL Server including login information, passwords, and server configurations.
14. In the **Username and Password** fields enter a user name and password as needed. If the VDP 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. In the space provided, 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 either:
 - 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 Items 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.*

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

Note: *When performing an Application Aware 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.*

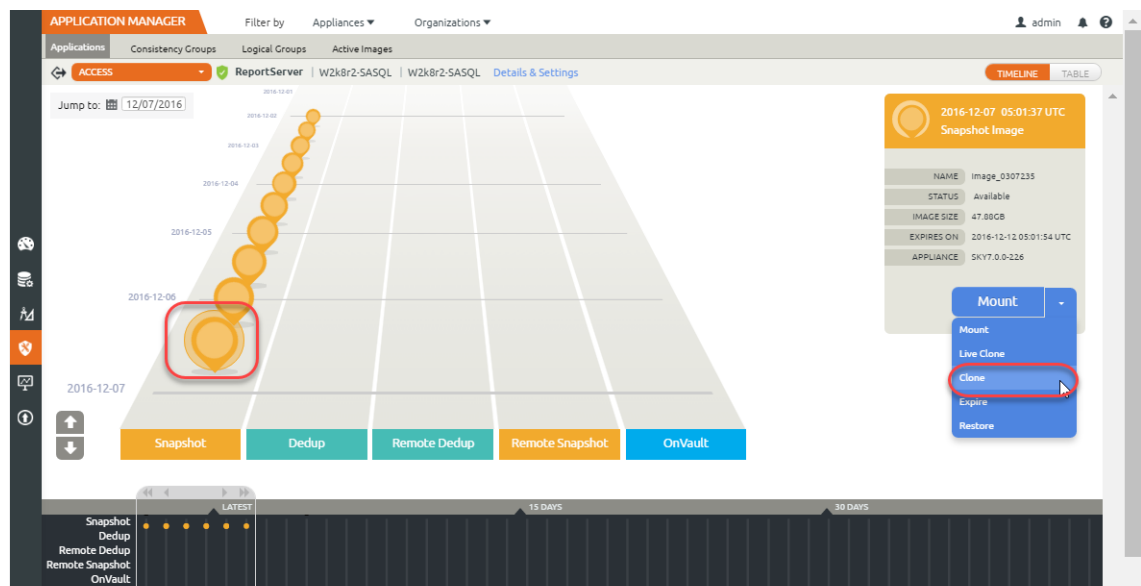
6 Cloning SQL Server Databases

You can clone (copying) a captured Microsoft SQL image to any host managed by your InfoSphere VDP Appliance. 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.

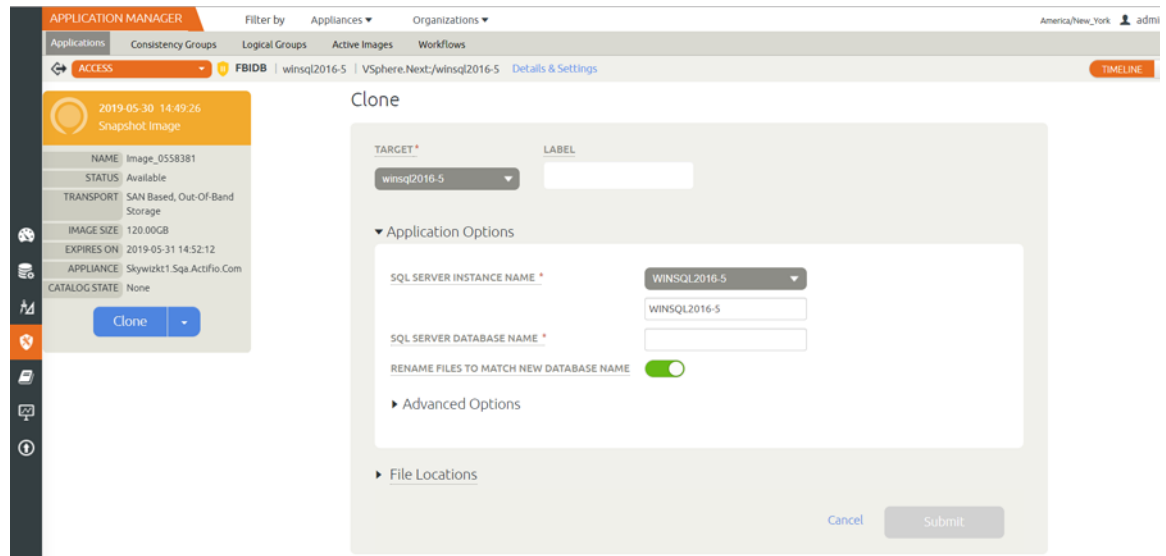
To clone a SQL database to a host:

1. From the left-hand navigation, click the Application Manager icon. The Applications page opens.
2. Select the application with the image that you want to clone, then choose Access from the drop-down list at the bottom right corner of the Applications page. The Access page opens listing captured images in the Timeline ramp view. Image types that support LiveClone creation include Snapshot and Remote Snapshot (StreamSnap images).

The Timeline ramp view is a time-based visualization of 7 days of captured images for the selected application. You can use your mouse scroll wheel or the up and down arrows in the bottom left corner of the page to move the timeline through the captured images and make a selection.



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



4. Select a host from the Select Host 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).
7. 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
8. Select a single volume or multiple volumes from Select Volumes To Clone. By default, all volumes are selected, and the first volume cannot be deselected.
9. Click **Submit**. You can see the cloned image in the Applications page (Applications > Managed Applications).

7 Mounting Encrypted SQL Data

InfoSphere VDP 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 33
- [Troubleshooting SQL Server Encryption](#) on page 35
- [SQL Server Master Key, Encryption Certificate, and Password Procedures](#) on page 36

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 an application aware 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

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 either perform a [Manual Query to Determine if SQL Encryption is Enabled](#) or use Microsoft's SQL Server Management Studio's user interface (SSMS):

Manual Query to Determine if SQL Encryption is Enabled

You can use a 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

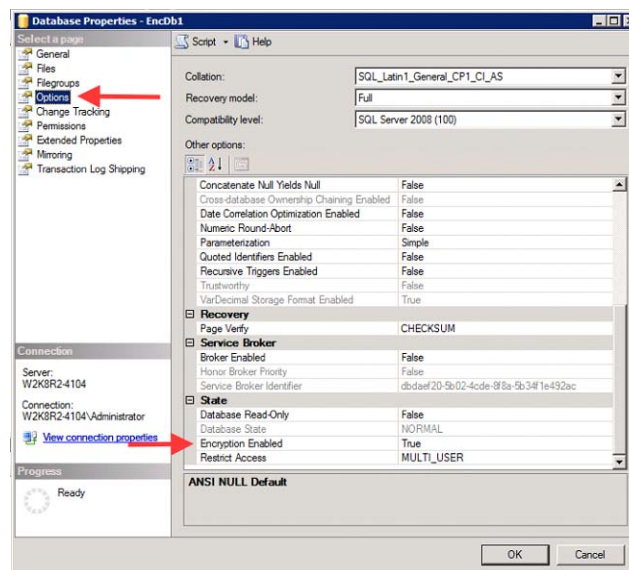
To use SSMS to determine if encryption is enabled on a database:

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:

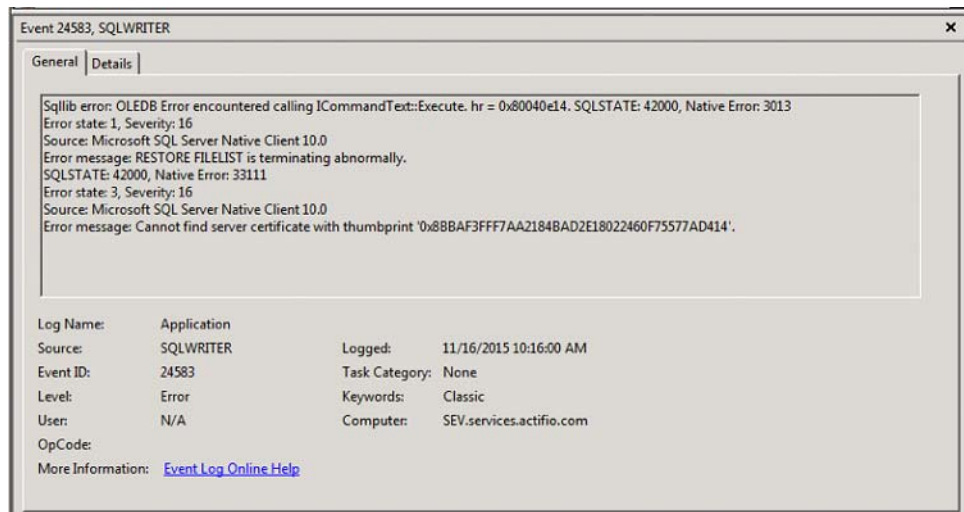


Database Properties

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**.

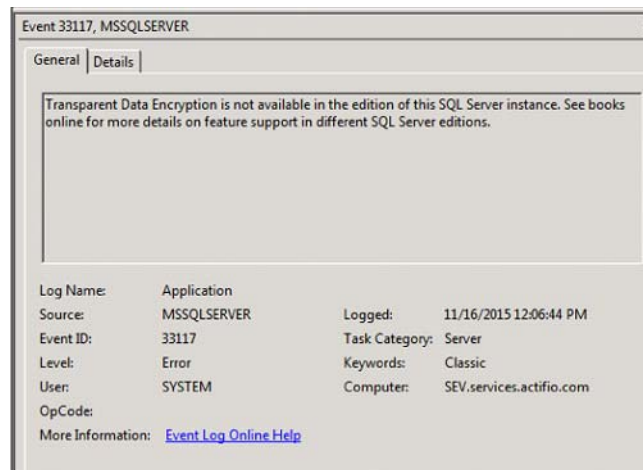
Troubleshooting SQL Server Encryption

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

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 InfoSphere VDP Appliances. The following procedures are provided as a convenience. If questions arise, consult Microsoft's detailed information on importing and exporting 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 a Source Password

If an SQL Server database on one SQL instance will be mounted to another SQL instance, you must manually make a copy of the to be mounted database's Server Master Key, encryption certificate, and password. Once the copy is made, you can 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 a location on 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
```


8 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. To mount the database as a virtual application see [Mounting an SQL Database as a Virtual Application](#) on page 34.

This chapter describes:

- [Microsoft SQL Server Database Restore Overview](#) on page 38
- [Restoring Microsoft SQL Instances and Databases](#) on page 39
- [Restoring a SQL Server Database to a Different Host](#) on page 40
- [Restoring SQL Server Databases in a Consistency Group](#) on page 40
- [Restoring SQL System Databases](#) on page 41
- [Restoring to an SQL Server Cluster](#) on page 42

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.

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.

The mounting process is wizard driven and varies slightly depending on whether you are mounting 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.
- The Restore operation cannot be performed *from* a remote InfoSphere VDP Appliance. However, you can restore with a remote-dedup image on the source InfoSphere VDP Appliance.
- Turn off the SLA options: **Run Schedule** and **Expire Data** for the application's policy template.
- Wait for running jobs to finish.

Microsoft SQL Server Database Restore Overview

The Restore function initiates all copy data recovery options for an InfoSphere VDP 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. Restore is the only data access operations that moves data. 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: IBM InfoSphere 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 VDP 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 IVGM by specifying Restore With Recovery, the SQL Server database will be restored and brought online without applying logs.

IVGM supports the following common use cases when restoring a 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 (Application Aware mount) when you encounter a corrupted SQL Server database: You can use an Application Aware 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.

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 6, 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 InfoSphere VDP Appliance. However, you can use a remote dedup image for a restore operation or a Clone operation on the source InfoSphere VDP Appliance.

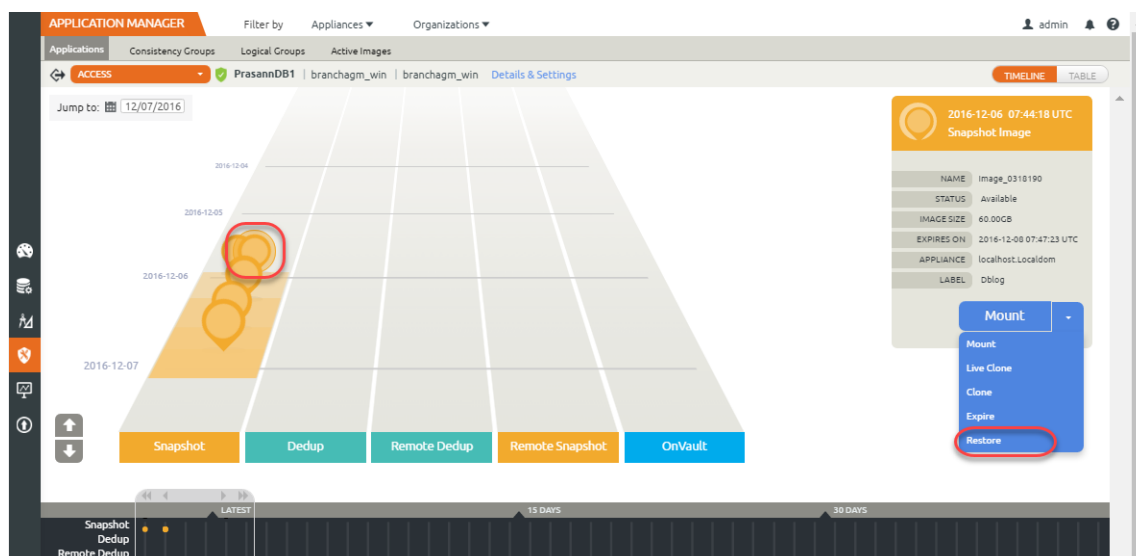
To run this procedure:

- 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 IVGM from a remote InfoSphere VDP Appliance. However, you can restore with a remote-dedup image on the source InfoSphere VDP 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:

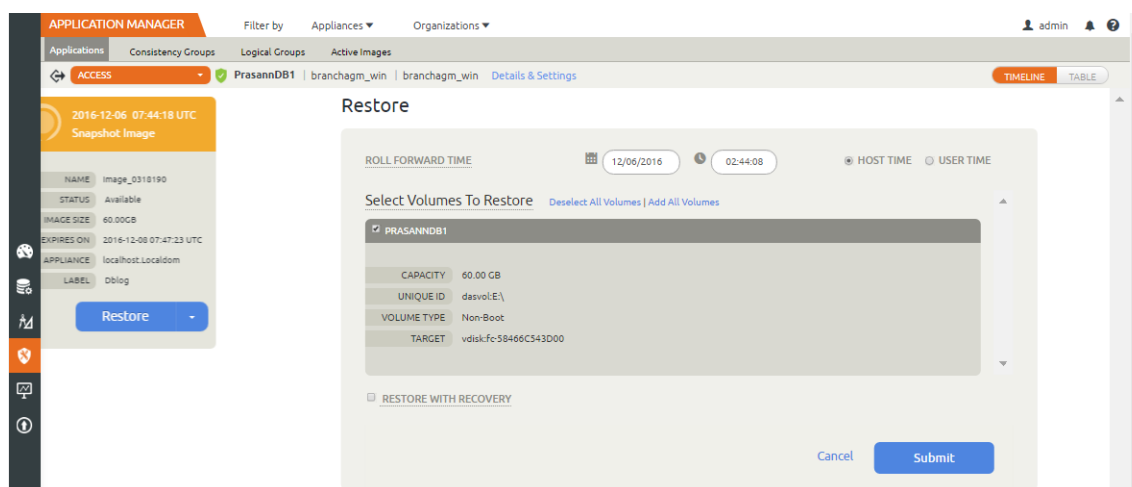
1. From the IVGM left-hand navigation, click the Application Manager icon. The Applications page opens.
2. Select the Microsoft SQL Server database that has the image that you plan to restore.
3. Choose **Manage SLA** from the drop-down list at the bottom right corner of the Applications page. The Manage SLA page opens. 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 Timeline ramp view is a visualization of seven days of captured images for the application. You can use your mouse scroll wheel or the up and down arrows in the bottom left corner of the page to move the timeline through the captured 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. If the selected database does not currently 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 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.
7. Select a single volume or multiple volumes to restore. By default all the volumes are selected.
8. Check the Restore With Recovery check box if you do not intend to roll the database logs. Restore with recovery brings the restored database on line. Once the database is online, logs cannot be applied.
9. 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 Microsoft SQL Server database image is restored, IVGM 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 because of corruption, or if the old database server is being replaced with a new server, then use a Clone operation as detailed in [Chapter 6, Cloning SQL Server Databases](#).

Restoring SQL Server Databases in a Consistency Group

Use caution when restoring Microsoft SQL Server databases in a consistency group. When you restore databases in a consistency group, all databases in the consistency group are overwritten. If you do not want to overwrite all databases in a consistency group, clone a single database: see [Chapter 6, Cloning SQL Server Databases](#).

For an SQL Server Failover instance, the database is always restored to the active node. The InfoSphere VDP Appliance 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

IVGM 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 39.

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.

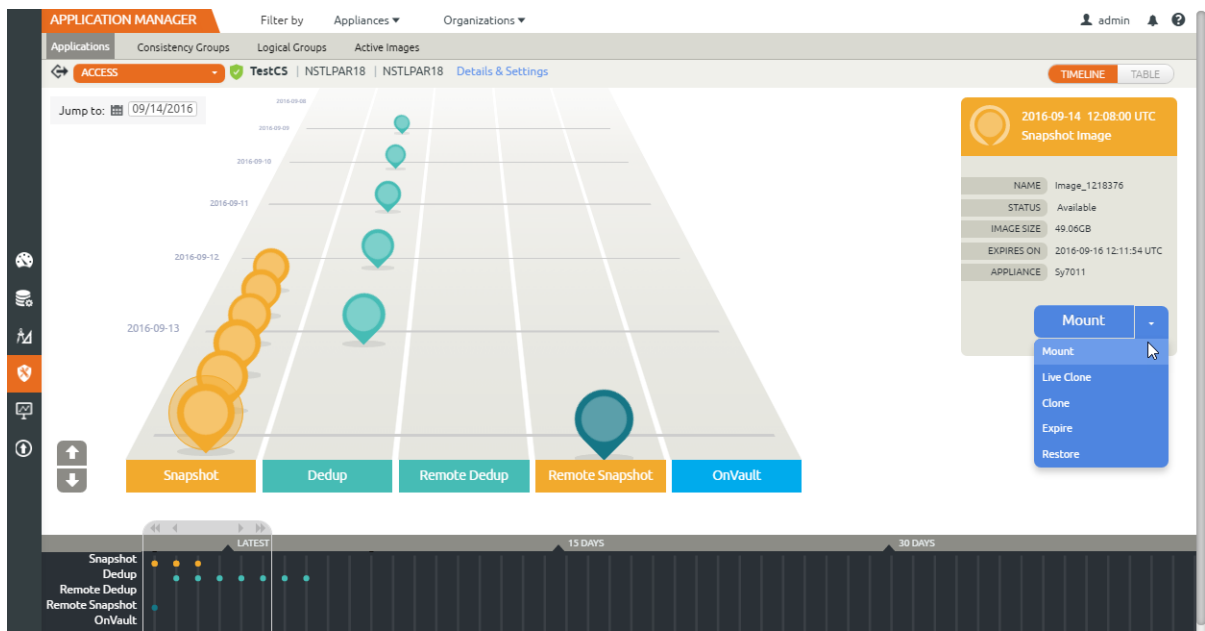
To manage the SQL system databases associated with an SQL user database, use a single policy template and resource profile to capture them in a consistency group. Doing so minimizes the consumption of system resources (VDisks) and reduces the number of jobs required to capture data.

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

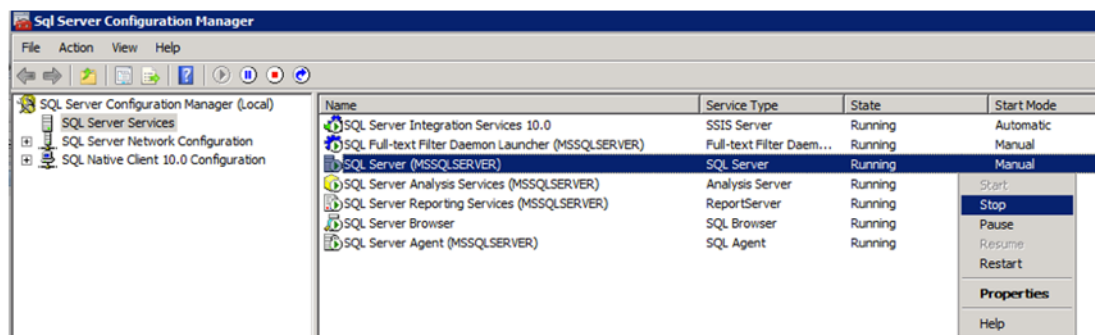
- master
- msdb
- model

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 Application Manager, 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. IBM InfoSphere 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 43.

9 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 43
2. [Restoring the Database on the Primary AAG Node](#) on page 43
3. [Synchronizing Secondary Databases to the Restored Primary Database](#) on page 44
4. [Rebuilding the SQL AlwaysOn Availability Group](#) on page 46

This chapter also provides information on error messages you can encounter during this process. See [Error Messages](#) on page 47 for details.

Before restoring, determine if the database was replicated to another site and if it was, was it replicated using IBM InfoSphere's 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 have to select a log that falls 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. You can do this by performing application aware mounts of the most recent images and checking them. You can mount multiple images simultaneously. Application aware mounts are detailed in [Chapter 5, Mounting an SQL Server Database as a New Virtual Database](#).

Restoring the Database on the Primary AAG Node

After you have identified the last known good version of the primary database, use IVGM to restore the primary production database from that image.

Restore operations are detailed in [Chapter 8, 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. Doing so will prevent you from applying logs to the secondary database(s) and will 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 AAAG 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 44.
- 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 IBM InfoSphere Clone image:
 - o Secondary databases that are captured by an IBM InfoSphere Production to Mirror policy and reside on a remote InfoSphere VDP Appliance: see [Restoring a Secondary SQL Server Database From an IBM InfoSphere Mirror Copy](#) on page 44.
 - o Secondary databases that are captured by an IBM InfoSphere Production to Dedup DR policy and reside on a remote InfoSphere VDP Appliance: see [Restoring a Secondary SQL Server Database From an IBM InfoSphere Dedup DR Copy](#) on page 45.

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:

- a. 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];
```
- b. Drop the secondary databases:

```
drop database [DATABASENAME];
```
- c. Join the replica database to AAG in **Full synchronization**. See [Rebuilding the SQL AlwaysOn Availability Group](#) on page 46 for details.

Restoring a Secondary SQL Server Database From an IBM InfoSphere Mirror Copy

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

You can restore a primary database from a second InfoSphere VDP 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:
 - o **StreamSnap** along with the database and are the exact logs to use for the restored primary.
 - o **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 IBM InfoSphere Dedup DR Copy](#) on page 45 for details.

To restore a secondary database from an IBM InfoSphere Mirror copy:

1. Log in to the InfoSphere VDP Appliance that manages the Mirror copy of the database.
2. Via the Application 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 6, 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 46.

Restoring a Secondary SQL Server Database From an IBM InfoSphere Dedup DR Copy

This approach to restoring a secondary database can be used if the primary database is captured by an InfoSphere VDP Appliance and then replicated to another InfoSphere VDP 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 IBM InfoSphere Mirror copy of the primary database is not available
- Dedup DR versions of the database reside on a second InfoSphere VDP 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 IBM InfoSphere Dedup DR copy:

1. Log in to the InfoSphere VDP 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 43.
3. Clone the database image as described in [Chapter 6, 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 46.

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:



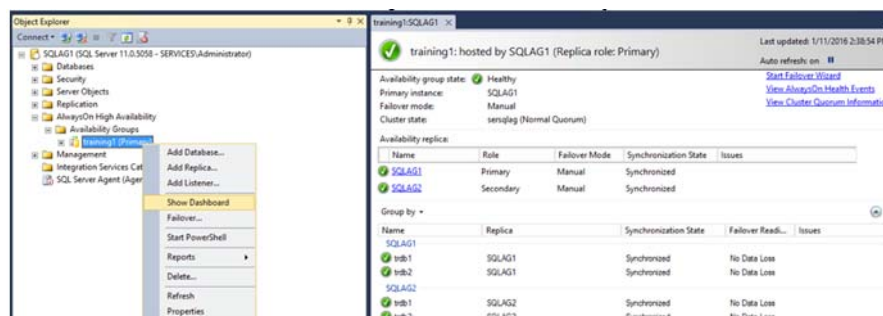
Select AAG

2. When prompted, select the primary database and click **Next**. You will be prompted to select a synchronization method:



Synchronization Method

3. 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 IBM InfoSphere Mirror Copy](#) on page 44 or [Restoring a Secondary SQL Server Database From an IBM InfoSphere Dedup DR Copy](#) on page 45. The Join only operation will form the primary and secondary databases back into an AAG.
4. Monitor the progress of the operation:

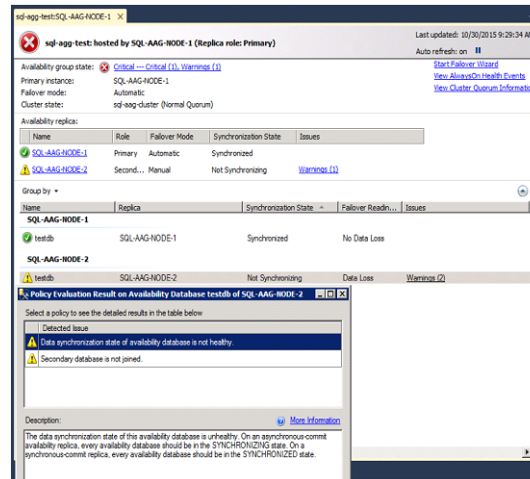


Monitor Operation

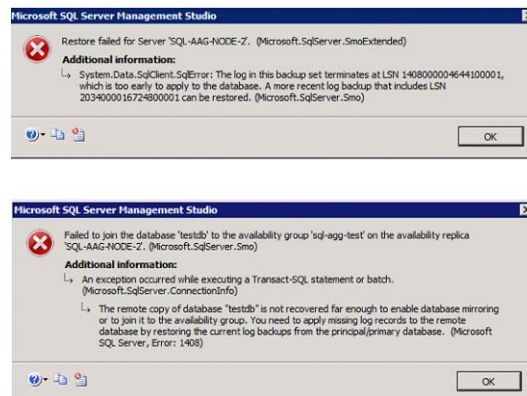
5. Once the operation is complete and both the primary and secondary databases are in sync, re-enable the IBM InfoSphere capture jobs for the AAG.

Error Messages

If you attempt to rebuild an AAG from an IBM InfoSphere Mirror copy that does not fall within the required range as specified in [Restoring a Secondary SQL Server Database From an IBM InfoSphere Mirror Copy](#) on page 44, you will encounter one or more of the following errors:

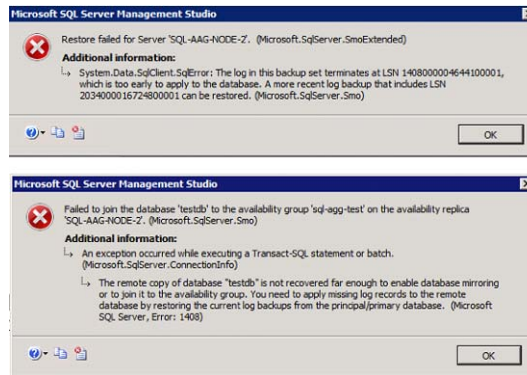


Error Example



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

10 Using Mounts to Create SQL AlwaysOn Availability Groups

This chapter details how to:

- [Creating an SQL Server AAG in an IBM InfoSphere Snapshot Pool on page 49](#)
- [Creating an SQL Server AAG Outside of An IBM InfoSphere Snapshot Pool on page 49](#)
- [Creating the New SQL Server AlwaysOn Availability Group on page 50](#)

Creating an SQL Server AAG in an IBM InfoSphere 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 IBM InfoSphere Snapshot Pool, it will also 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 Database as a Virtual Application](#) on page 34.
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 50.


Creating an SQL Server AAG Outside of An IBM InfoSphere Snapshot Pool

To create an AAG that resides outside of an IBM InfoSphere Snapshot pool, as described in [Chapter 6, Cloning SQL Server Databases](#).

Creating the New SQL Server AlwaysOn Availability Group

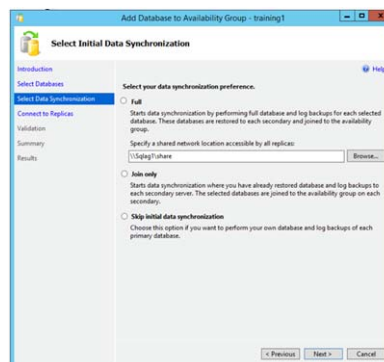
After all databases for the new AAG have been recovered, use SQL Studio or T-SQL to create the new AAG:

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:



Synchronization Method

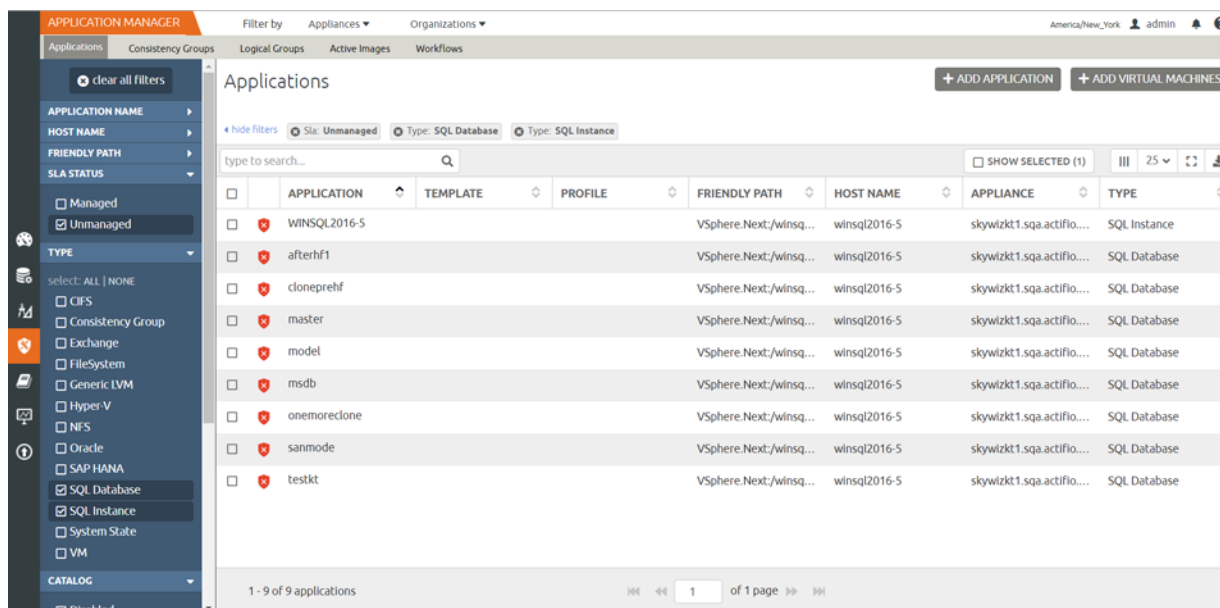
4. Select **Join only**: The Join only operation will form the primary and secondary databases into an AAG.
5. Monitor the progress of the operation:



Monitor Progress

When the operation is complete and both the primary and secondary databases are in sync.

6. If it is not already present, install the appropriate VDP Connector on AAG's host.
7. If auto discover applications is not enabled, manually discover applications on the AAG host.
8. After discovery is complete, the host's AAGs will appear in the Application Manager:



Discovered AAGs

9. Capture members of an AAG just like any other database. See [Protecting Microsoft SQL Server Instances and Databases](#) on page 15 for details.

