
Creating Automated Workflows for SQL Server Databases



Copyright, Trademarks, and other Legal Matter

Copyright © 2009 - 2020 Actifio, Inc. All rights reserved.

Actifio[®], AnyIT[®], Dedup Async[®], OnVault[®], Enterprise Data-as-a-Service[®], FlashScan[®], AppFlash DEVOPS Platform[®], Copy Data Cloud[®], and VDP[®] are registered trademarks of Actifio, Inc.

Actifio Sky[™], Actifio One[™], and Virtual Data Pipeline[™] are trademarks of Actifio, Inc.

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

Actifio, Inc., is a provider of data protection and availability products. Actifio's technology is used in products sold by the company and products and services sold and offered by its commercial partners. The current list of Actifio patents is available online at: <http://www.actifio.com/patents/>

Actifio believes the information in this publication is accurate as of its publication date. Actifio reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." ACTIFIO, INC. MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

This software and the associated documentation are proprietary and confidential to Actifio. Use, copying, and distribution of any Actifio software described in this publication requires an applicable software license. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

Actifio strives to produce quality documentation and welcomes your feedback. Please send comments and suggestions to docs@actifio.com.

Contents

Chapter 1 – Introduction to Provisioning Environments With Workflows	1
Workflow Uses.....	2
Workflow Benefits.....	3
Capturing Data for Workflows	3
Workflow Data Access Methods.....	4
Chapter 2 – Workflow Access Control	5
Configuring Roles, Organizations, and Users.....	5
Example Role for Limited Workflow Access.....	6
Example Organization for Limited Workflow Access.....	7
Example User for Limited Workflow Access.....	8
Chapter 3 – Defining Workflows	9
LiveClone Workflow: Scrubbed Virtual Application	10
LiveClone Workflow: Scrubbed Application Data.....	16
Direct Mount Workflow: Virtual Application.....	20
Direct Mount Workflow: Application Data.....	24
Chapter 4 – Running Workflows On Demand	27
LiveClone Run Now: New Scrubbed Virtual Application.....	28
LiveClone Run Now: Re-Provisioning a Scrubbed Virtual Application.....	31
LiveClone Run Now: Scrubbed Application Data	34
Direct Mount Run Now: Virtual Application.....	36
Direct Mount Run Now: Application Data	39
Chapter 5 – Workflow Pre and Post Scripts	41
Introduction	41
Environment Variables	42
High-Level Examples	43
Provision a New Database, Log Capture Not Enabled.....	43
Refresh a Previously Provisioned Database, Log Capture Not Enabled.....	43
Provision a New Database with Log Capture Enabled.....	43
Variables Used to Mount Databases	44

Variables Set for a Database Mount Without Logs.....	44
Variables Set for a Database Reprovision With Logs.....	44
Example Script.....	49

1 Introduction to Provisioning Environments With Workflows

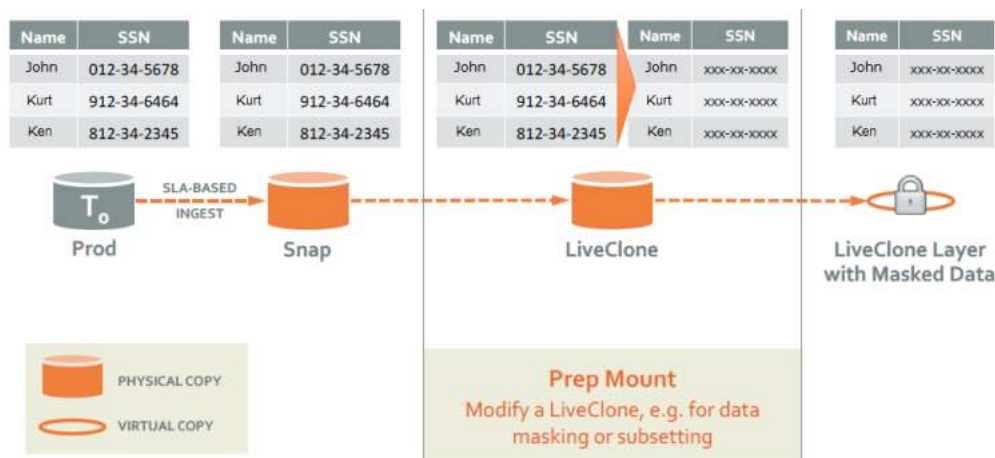
While Actifio SLA Policy Templates govern the automated capture of production data, Actifio Workflows automate access to captured data. Workflows can run according to a schedule or on demand.

Workflows are built with captured production data. Workflows can present the captured data as a LiveClone, a virtual application, or as just the application data.

For those applications that contain sensitive data, a Workflow can include a step that creates a LiveClone and then automatically runs a script to mask the sensitive data.

The LiveClone with the masked data can then be mounted as a virtual application or the LiveClone's data to provision a work environment. For example, the following figure illustrates a Workflow that:

1. Captures production data that contains Social Security numbers
2. Prep-mounts the captured data as a LiveClone so a script can scrub the Social Security numbers from the data
3. Mounts the scrubbed LiveClone as a virtual application via an Application Aware mount.
4. Flags the scrubbed virtual application as non-sensitive.



Once a Workflow completes, users can access the server on which the virtual application has been mounted and use the data. Access to the data is controlled by the server on which the data is mounted.

This chapter provides an overview of Actifio Workflows, including:

- [Workflow Uses](#) on page 2
- [Workflow Benefits](#) on page 3
- [Capturing Data for Workflows](#) on page 3
- [Workflow Data Access Methods](#) on page 4

Workflow Uses

Some of the more common uses for Workflows include:

Provisioning Work Environments

Teams such as test and development can take advantage of a Workflow's ability to present captured production data as a virtual application. When production data contains sensitive information, the data can be scrubbed prior to making it available to users.

The ability to limit access to specific features and functions allows users to provision and re-provision their environments without putting production data at risk and without the assistance of database administrators or system administrators.

Users who have access to a virtual application provisioned by a Workflow can also capture that virtual application. The new application can be worked on and tested separately from the regularly updated data provided by the Workflow. For example, by peeling off the data as a new application, tests such as user acceptance tests can be performed on the new application. Tested applications can be updated and retested until they are ready for release.

A test and development Workflow could consist of:

1. Defining a schedule for updating the LiveClone with the latest production data. Scheduled Workflows can also be run on demand.

Note: You can run a Scheduled Workflow on-demand at any time. You cannot convert an On-Demand Workflow to a Scheduled Workflow after it has been created.

2. Generating a new, or updating an existing LiveClone from a selected production image marked as sensitive.
3. Employing scripts to scrub the LiveClone of sensitive data.
4. Marking the scrubbed application data as non-sensitive.
5. Mounting the scrubbed image as a virtual application (Application Aware Mount) to a specified host.
6. Capturing the scrubbed virtual application as a new application.

For an example of this type of Workflow, see [LiveClone Workflow: Scrubbed Virtual Application](#) on page 10.

Reporting

To analyze and report on production data without interfering with the production environment, you can define a Workflow that generates a virtual application and then executes a script that analyzes and reports on the virtual application. Once the report is generated, the mounted data can be unmounted and deleted.

A reporting Workflow could consist of:

1. Mounting data as a virtual application
2. Executing a script that analyzes the virtual application and outputs the results in a format that can be read by a third party reporting tool.
3. When the script is complete, unmounting and deleting the mounted image.

For an example of this type of Workflow, see [Direct Mount Workflow: Virtual Application](#) on page 20.

Data Warehousing

Teams can automate data warehousing, especially when data must be extracted, transformed, and loaded (ETL). A data warehousing Workflow could consist of:

1. Mounting data as a virtual application.

2. Executing a script that extracts, transforms, and loads the data.
3. When ETL is complete, unmounting and deleting the mounted image.

For an example of this type of Workflow, see [Direct Mount Workflow: Virtual Application](#) on page 20.

Copying to Removable Media

Workflows can automate copying selected data to a removable media such as tape for long-term retention. A Workflow that makes application data available for copying to removable media could consist of:

- Mounting just the application data.
- Executing a script that will initiate a third party application to copy the data to removable media.
- When the copy is complete, unmounting and deleting the mounted image.

For an example of this type of Workflow, see [Direct Mount Workflow: Application Data](#) on page 24.

Workflow Benefits

The Actifio appliance allows you to capture production data via a schedule or on demand. Once the data is captured it can be used in a Workflow to provision work environments.

With Actifio Workflows:

- Work environments can be updated with production data as soon as new data is available without having to wait days or even weeks for a DBA or system administrator to do the provisioning.
- Work environments can be updated on a predefined schedule or on demand.
- User Acceptance Testing (UAT) is streamlined. Actifio can provision virtual applications (virtual copies of production applications) in UAT environments using the same protocols and environments that exist in production.
- DBAs and system administrators are freed from constantly provisioning and refreshing work environments.
- Users can quickly and efficiently reproduce and address issues using the latest production data in a secure, isolated environment.
- Software updates can be applied and tested, using the latest production data before being released into a production environment.
- Multiple teams can work in parallel and instantly access their own copy or virtual copy of production data.

Actifio Workflows take provisioning one step further by providing the ability to scrub images of sensitive data. This ability allows users to self-provision their environments with virtual applications of production data while at the same time maintaining data security.

Capturing Data for Workflows

Workflows are made up of captured application data. Data captured for other purposes can be used by an Actifio Workflow. If you are capturing data specifically for Workflow purposes, the capture operation is no different than capturing data for other purposes. From the Actifio Desktop:

1. Add servers that host application data.
2. Discover applications on the added servers.
3. Define Actifio Policy Templates and Resource Profiles (SLAs) according to your RPOs and RTOs.
4. Apply Actifio Policy Templates and Resource Profiles to applications.

For details on capturing application data, see the AGM online help.

Workflow Data Access Methods

Workflows present captured data as Application Aware Mounts, LiveClones, or Standard Mounts.

Application Aware Mount

The Actifio Application Aware mount function provides instant access to an application as a virtual application without actually moving data. Applications can be mounted on any application server. If the application is a database, and its logs have been captured, you can use the logs to roll the database forward to a desired point in time. Application Aware mounts are performed from the Actifio appliance and do not require manual intervention by application, database, server, or storage administrators.

Once an Application Aware Mount finishes, the virtual application can be captured like any other application. This capability is particularly useful when a problem is encountered with a virtual application. Users can roll back their copy of the virtual application to a point where the issue does not exist, while other users troubleshoot the problematic version.

LiveClone

A LiveClone is an independent copy of an application that can be refreshed on a schedule or on demand. A LiveClone can be scrubbed for sensitive data before being made available to users. This allows teams such as development and test to use production quality data without exposing sensitive data and interfering with the production environment. LiveClones are updated on demand or according to a schedule defined in the Workflow.

Direct Mount

Direct mounts present only an application's data. Direct mounts are updated on demand or according to a schedule defined in the Workflow.

2 Workflow Access Control

Actifio appliances come with a predefined Actifio Admin user. The Actifio Admin is a super user who has full control of and access to all features, functions, and resources of an Actifio appliance.

The Admin User can create users with various degrees of access. When creating Workflows, consider which users will have access to Workflows and what access they will have to an Actifio appliance's features, functions, and managed data - including access to sensitive data.

This chapter contains:

- [Configuring Roles, Organizations, and Users](#) on page 5
- [Example Role for Limited Workflow Access](#) on page 6
- [Example Organization for Limited Workflow Access](#) on page 7
- [Example User for Limited Workflow Access](#) on page 8

Configuring Roles, Organizations, and Users

Ultimately the users created and their degree of access to an Actifio appliance and the data it captures are driven by your corporate policies and the needs of your organizations. This section describes one of the more common approaches to creating users and assigning them the roles and organizations needed to run and view Workflows.

Although any user can be granted privileges to create and run a Workflow, team leaders are often tasked with running Workflows on demand and provisioning the work environment for their team. Team leaders typically are not the users who create Workflows, but are often responsible for running Workflows.

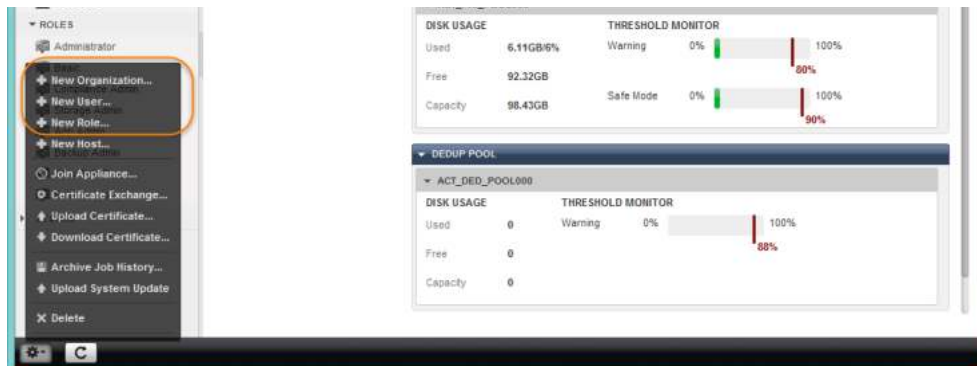
Giving team leaders the ability to run Workflows on demand frees DBAs and system administrators from having to update work environments.

A Workflow's on-demand capabilities allow team leaders to update their environments with refreshed production data as needed. For example, if a Workflow is scheduled to run every Sunday night, a team leader can run a Workflow on demand on a Wednesday and update the work environment with the latest production data.

To create an Actifio user who has just enough access to an Actifio appliance to run Workflows and to provision their environment with the latest Actifio Workflow updates, you can create:

- A role that has the necessary privileges to log in to the Actifio appliance, access the Application Manager, manage mounts, run Workflows, and view the status of a Workflow's progress.
- An organization that limits the user to Workflow data captured from specified application servers and if necessary, specific applications.
- A user who may or may not have access to sensitive data, and is assigned to an appropriate role and organization.

New roles, organizations, and users are created in the Domain Manager service. They are created by clicking the service menu icon on the bottom left-hand side of the Domain Manager service:



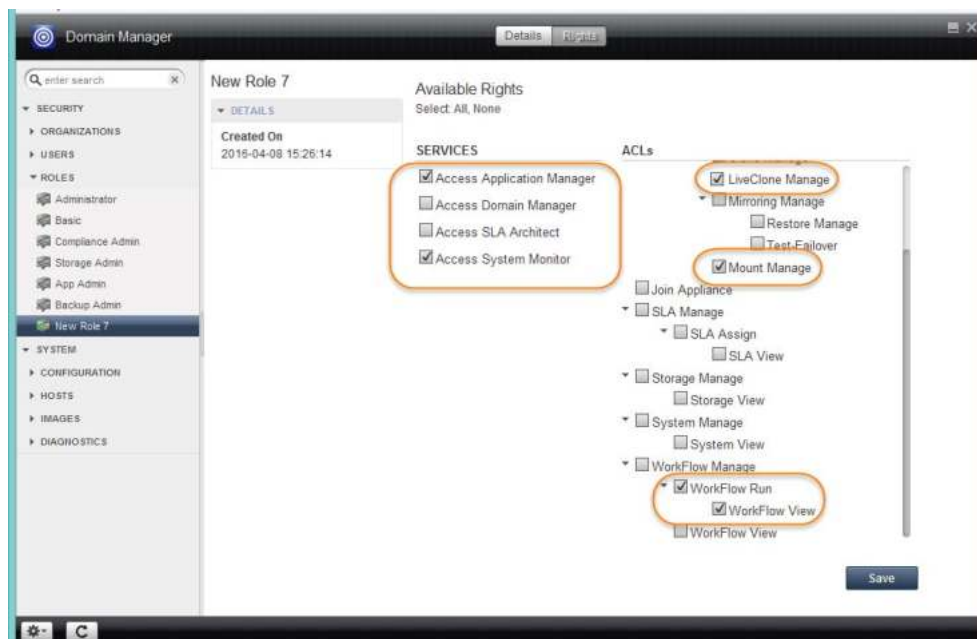
Details on creating new roles, organizations, and users can be found in the AGM online help.

Use the following examples as a quick reference for creating users who do not have access to sensitive data, and are assigned roles and organizations in such a way that they can only use an Actifio appliance to provision their work environment.

Example Role for Limited Workflow Access

The following screen capture is an example of a user role that limits a user to accessing an Actifio appliance's Application Manager, Manage LiveClones, Manage Mounts, and to view and run Workflows. It also allows that user to monitor the progress of a Workflow via the Actifio System Monitor. A single role can be assigned to multiple users.

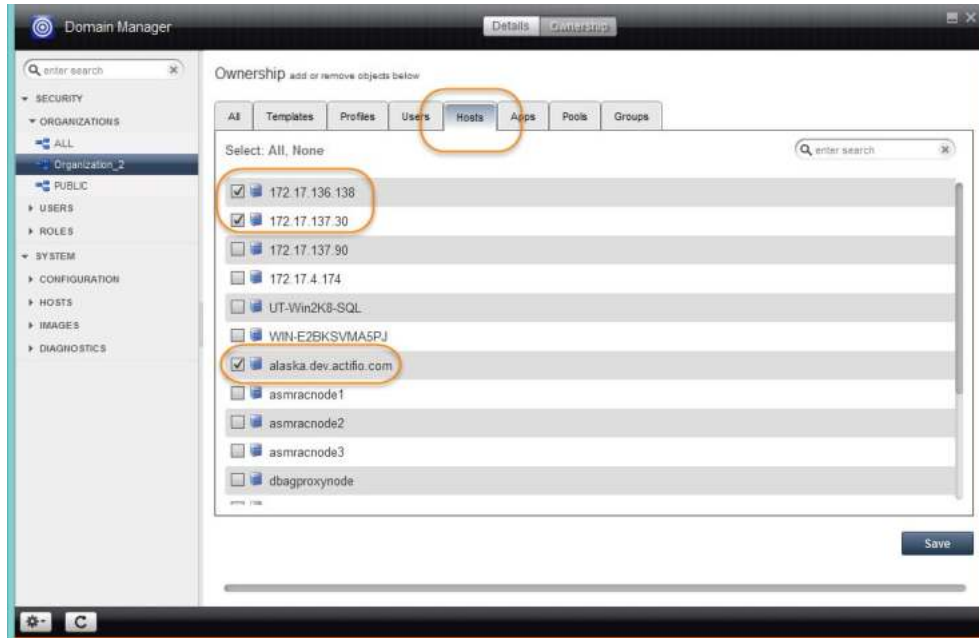
Note: Because a LiveClone in a LiveClone Workflow typically contains sensitive information, users assigned a role with Manage LiveClone rights must also have access to sensitive data. See [Example User for Limited Workflow Access on page 8](#).



Example Role for Limited Workflow Access

Example Organization for Limited Workflow Access

The following screen capture is an example of an organization that limits a Workflow user to specific application servers. If necessary, the organization could limit a Workflow user to specific applications on the servers. A single organization can be assigned to multiple users.



Example Organization for Limited Workflow Access

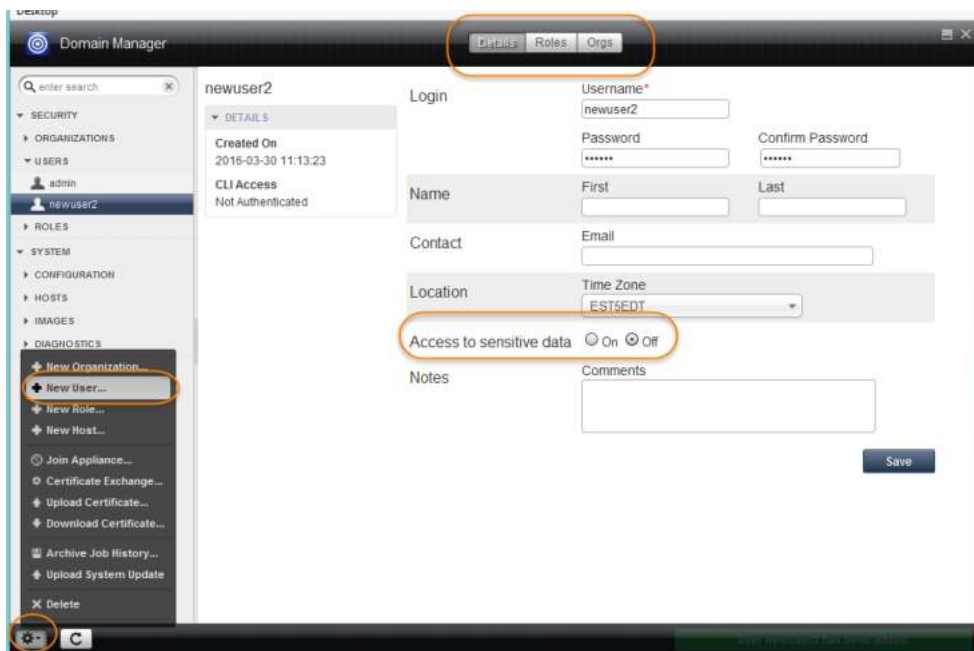
Example User for Limited Workflow Access

The following screen capture is an example of a new user who does not have access to sensitive data.

Note: Such users do not have access to sensitive data in a Workflow even if their assigned organizations contain sensitive data.

Roles and Organizations are assigned to a user by clicking the tabs at the top of the page. A single user can be assigned multiple roles and multiple organizations.

Note: Because a LiveClone in a LiveClone Workflow typically contains sensitive information, users assigned a role with Manage LiveClone rights must also have access to sensitive data. See [Example Role for Limited Workflow Access on page 6](#).



Example User for Limited Workflow Access

3 Defining Workflows

Workflows are defined in the Actifio Application Manager. Workflows use captured production data as the source data from which they generate:

LiveClones: LiveClones are typically used in Workflows when an application contains sensitive data which must be scrubbed before it is mounted and made available to users. The application data can be mounted with a standard mount as just the data, as a virtual application, or as both the data and as a virtual application. Updates to the LiveClone, scrubbing, and mounting can be done automatically via a schedule or on demand.

Direct Mounts: Direct Mounts are used when application data does not need to be scrubbed before it is mounted and made available to users. The application data can be mounted with a standard mount as just the data, as a virtual application, or as both the data and as a virtual application. Direct mounts make captured data available almost instantly without actually moving data.

Note: Workflows consist of a number of options. In most cases, when creating a Workflow, not specifying an option will allow a user to specify that option when running the Workflow on demand.

You can create an Actifio Workflow from captured SQL:

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

Options presented while creating a Workflow will vary according to the source that is selected. For example, databases on VMware VMs will have a **Map to all ESX hosts** option. If a LiveClone or virtual application is being made from a selected Consistency Group, you will have the opportunity to provide a name for the new LiveClone Consistency Group or virtual application Consistency Group.

This chapter provides step-by-step examples for defining a:

- [LiveClone Workflow: Scrubbed Virtual Application](#) on page 10
- [LiveClone Workflow: Scrubbed Application Data](#) on page 16
- [Direct Mount Workflow: Virtual Application](#) on page 20
- [Direct Mount Workflow: Application Data](#) on page 24

LiveClone Workflow: Scrubbed Virtual Application

LiveClone Workflows are typically used when the production data to be used contains sensitive information. The LiveClone Workflow allows you to define a LiveClone from production data and scrub the data before mounting it.

In the following example, a Workflow is defined that:

- Generates a new, or updates an existing LiveClone from a selected production image marked as sensitive.
- Defines a schedule for updating the LiveClone with the latest production data. Scheduled Workflows can also be run on demand.

Note: You can run a Scheduled Workflow on-demand at any time. You cannot convert an On-Demand Workflow to a Scheduled Workflow after it has been created.

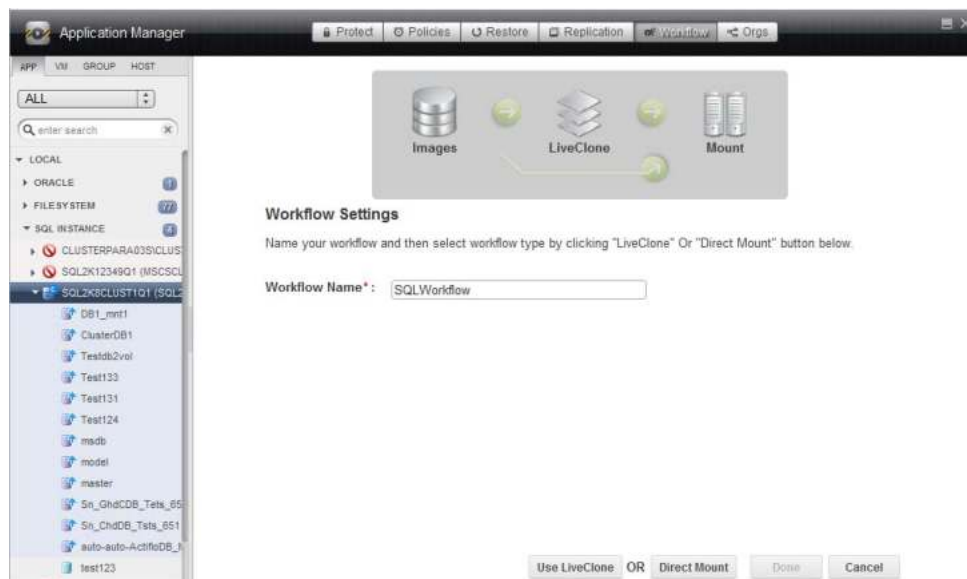
- Employs scripts to scrub the LiveClone of sensitive data.
- Marks the scrubbed virtual application as non-sensitive.
- Mounts the scrubbed image as a virtual application (Application Aware Mount) to a specified host.

Users with the proper access rights can run this Workflow as described in [LiveClone Run Now: New Scrubbed Virtual Application](#) on page 28 and [LiveClone Run Now: Re-Provisioning a Scrubbed Virtual Application](#) on page 31.

To define this Workflow:

1. Open the Actifio Application Manager.
2. In the Navigation pane, select the application to be scrubbed of its sensitive data.
3. Click the **Workflow** tab and the Workflow Settings page is displayed:

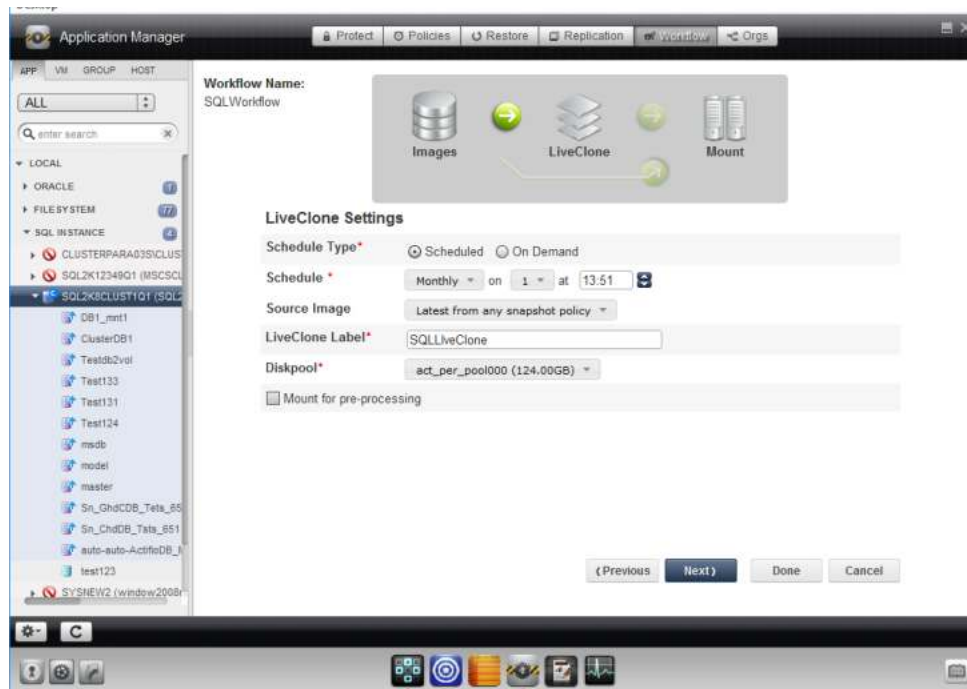
Note: If the application already has a Workflow, click the **Add** button.



New Workflow Settings Page

4. Specify a name for the Workflow. Workflow names cannot include special characters.

- Click **Use LiveClone** and the LiveClone Settings page is displayed:



LiveClone Settings Options

- For Schedule Type, select either **Scheduled** or **On Demand**. Scheduled Workflows can also be run on demand.

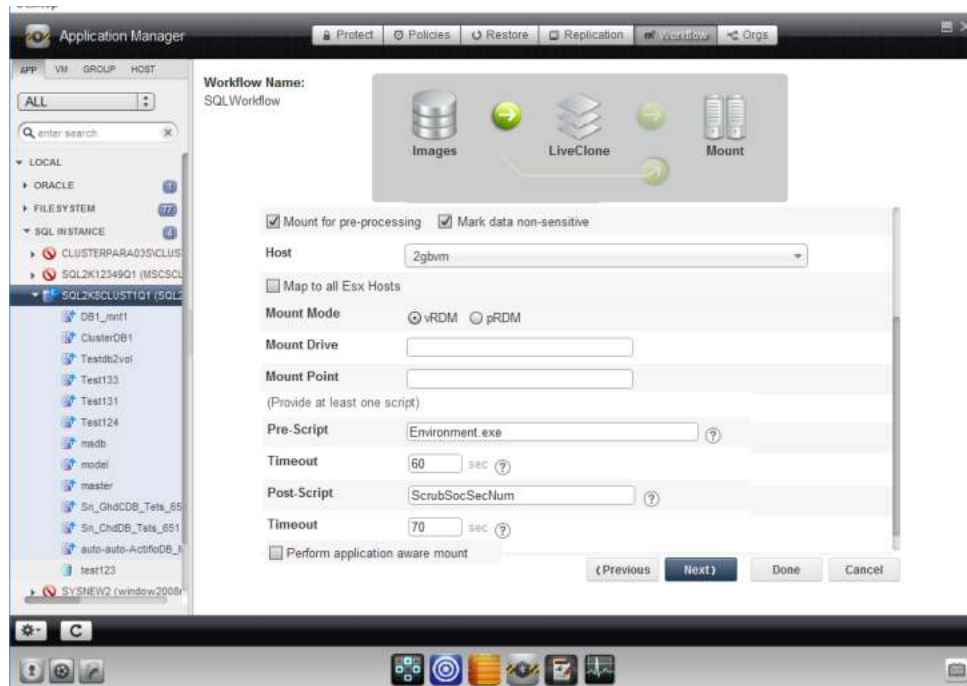
Note: You can run a Scheduled Workflow on-demand at any time. You cannot convert an On-Demand Workflow to a Scheduled Workflow after it has been created.

- If you selected **Scheduled**, define a schedule to run the Workflow. The time selector uses a 24-hour clock.
- For **Source Image**, select the latest snapshot image from the policy that produces the image to scrub.

Note: If you selected On Demand, the user running the Workflow can select the source snapshot image at run time.

- Give the LiveClone a label to make it easily identifiable.
- Select the Snapshot Pool in which the LiveClone will reside. Ensure enough storage space is allocated to accommodate the LiveClone. If changes are necessary, see the AGM online help for details on configuring Snapshot Pools.

Check **Mount for pre-processing** and the following screen provides the opportunity to mount the LiveClone to a host and then scrub the mounted image for sensitive data:



LiveClone Pre-Processing Options

Note: You may have to scroll down to see all options.

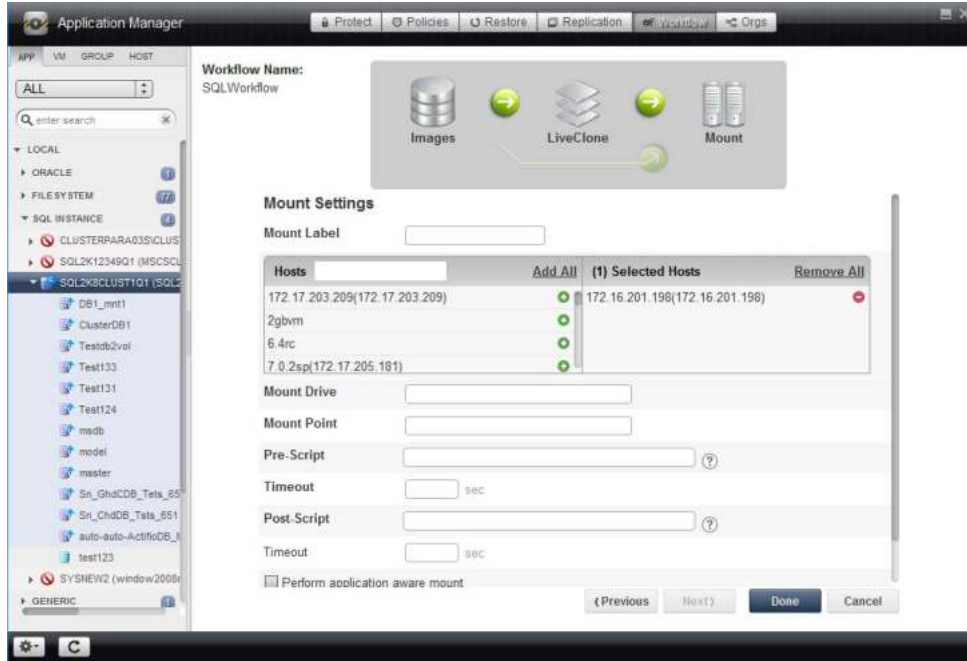
11. The **Mark data as non-sensitive** option only appears if the selected application is marked as containing sensitive data. Check this option if the application will not contain sensitive data after scrubbing. This will allow users who do not have access to sensitive data to access the Workflow and provision their work environments. See [Example User for Limited Workflow Access](#) on page 8 for details.
12. Select a host on which the LiveClone will be mounted.
13. If the database is on a VMware VM, you will be given the option to ensure connectivity by checking **Map to all ESX hosts**.
14. Add a **Mount Drive** (Windows) if you want the image mounted to a location other than the default. If you do not specify a mount point, then the Actifio Connector will assign a mount location.
15. Enter a timeout that will allow enough time for the script to run.
16. Specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting or unmounting the application. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
17. Enter a timeout that will allow enough time for the script to run.
18. Specify a **Post Script** as needed. The post script in this example will be used to scrub the application of sensitive information after it has been mounted. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
19. Check **Perform application aware mount** and the following application specific options appear:



LiveClone Application Specific Options

Note: If you have selected a consistency group, in the space provided enter a name for the LiveClone consistency group.

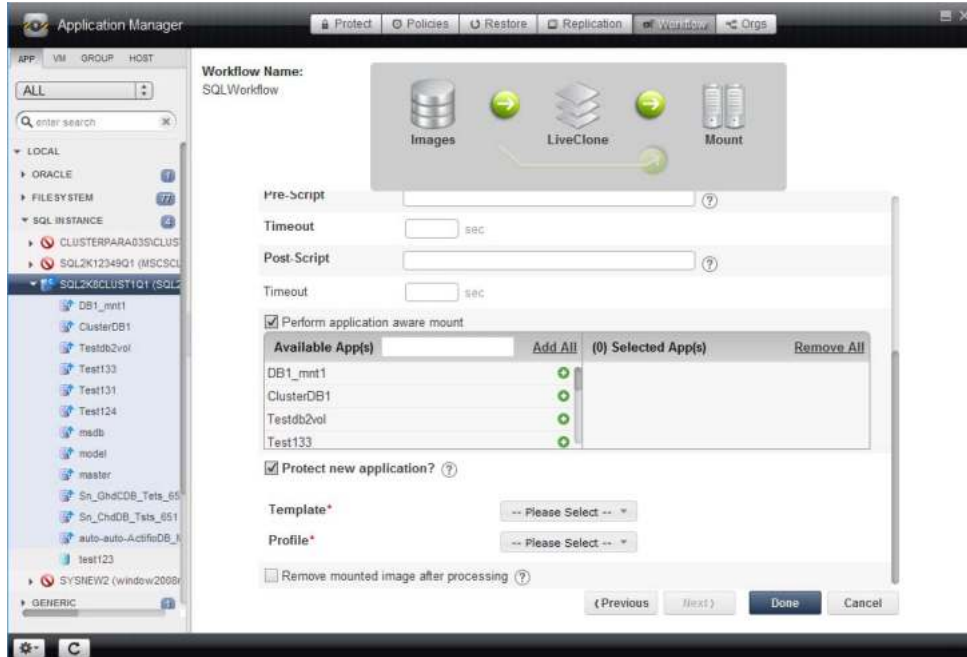
20. Enter an **SQL Server Instance Name** for the LiveClone instance.
21. Enter a prefix that will be appended to the newly provisioned database name.
22. Check **Recover Database After Restore**. Doing so leaves the database in a state where if logs are available they can be applied to bring the database to a specific point in time.
23. Check **Recover User Logins**. Doing so back ups the master database which records all the system-level information for a SQL Server. This includes login information, passwords, and server configurations.
24. In the spaces provided, enter a user name and password. The specified user must have the SQL Server Database role(s) necessary to perform all of the operations specified in the Workflow. For details on Microsoft SQL Server roles see the AGM online help.
25. Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand and provision and re-provision any server to which they have access with the virtual application.
26. Click **Next** and the Mount Settings options are displayed. Use these options to limit what a user running this Workflow on demand can specify.



LiveClone Mount Settings

27. Specify a **Mount label** to make the mount easily identifiable.
28. Click the green + icon next to a server name to specify a server on which the virtual application will be automatically mounted. Once set, users running the Workflow will not be able to specify another server.
29. Specify a **Mount Mode**, **Mount Drive**, or **Mount Point** as needed. Once specified, users running the Workflow will not be able to specify these mount settings.
30. **Pre Script** is a script used to prepare the environment prior to mounting or unmounting the virtual application. Do not confuse this script with the script you specified to scrub data. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
31. Enter a timeout that will allow enough time for the script to run.
32. **Post Script** is a script used to act on the virtual application after it has been mounted or after it is unmounted. Do not confuse this script with the script you specified to scrub data. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
33. Enter a timeout that will allow enough time for the script to run.

Note: Scripts are usually used along with the **Remove mounted image after processing** option.



LiveClone Mount Settings (Continued)

Note: You may have to scroll down to see all options.

34. Specify whether to protect the virtual application defined by this Workflow as a new application.
35. If you are protecting as a new application, select a template and profile to protect the new application.
36. **Remove mounted image after processing.** This is typically checked when you employ a script to process mounted data. Once the script finishes its task this option will unmount and delete the virtual application. See [Reporting](#) on page 2, [Data Warehousing](#) on page 2, and [Copying to Removable Media](#) on page 3.
37. Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand and provision and re-provision the server as specified in this section.

LiveClone Workflow: Scrubbed Application Data

LiveClone Workflows are typically used when the production data to be used contains sensitive information. The LiveClone Workflow allows you to define a LiveClone from production data and then scrub the LiveClone's data before mounting it.

In the following example, a Workflow is defined that:

- Generates a new, or updates an existing LiveClone from a selected production image that is marked as sensitive.
- Defines a schedule for updating the LiveClone with the latest production data. Scheduled Workflows can also be run on demand.
- Employs scripts to scrub the LiveClone of sensitive data.
- Marks the mounted, scrubbed data as non-sensitive.
- Mounts the scrubbed data (standard mount) to a specified host.

Users with the proper access rights can run this Workflow as described in [LiveClone Run Now: Scrubbed Application Data](#) on page 34.

To define this Workflow:

1. Open the Actifio Application Manager.
2. In the Navigation pane, select the application to be scrubbed of its sensitive data.
3. Click the **Workflow** tab and the Workflow Settings page is displayed:

Note: If the application already has a Workflow, click the **Add** button.



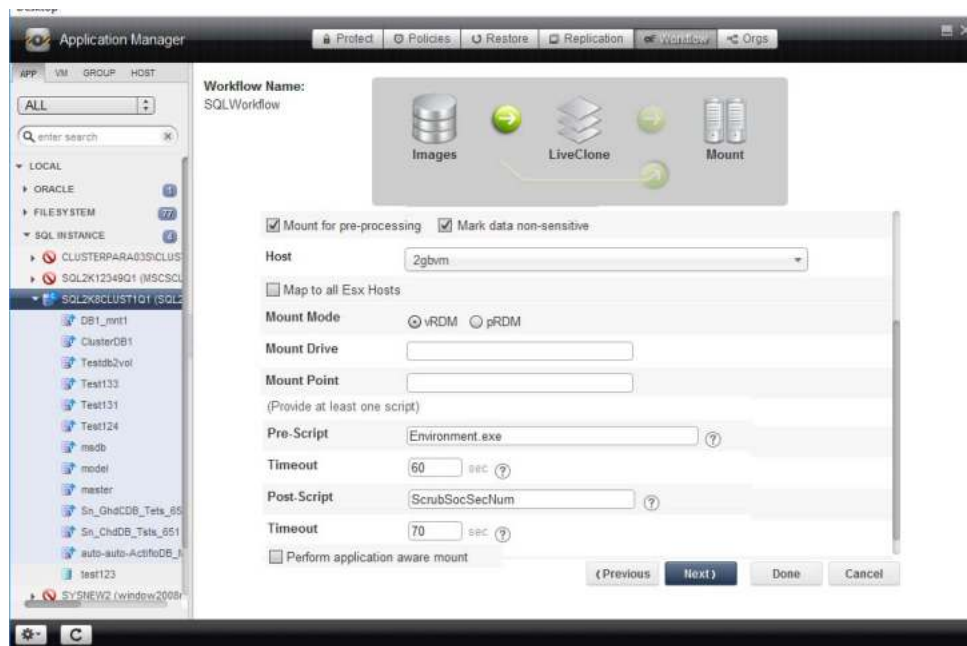
New Workflow Settings Page

4. Specify a name for the Workflow. Workflow names cannot include special characters.
5. Click **Use LiveClone** and the LiveClone Settings page is displayed:



LiveClone Settings Options

6. For Schedule Type, select **Scheduled**. Scheduled Workflows can also be run on demand.
7. In **Schedule**, define a schedule to run the Workflow. The time selector uses a 24-hour clock.
8. For **Source Image**, select the latest snapshot image from the policy that produces the image.
9. Give the LiveClone a label to make it easily identifiable.
10. Select the Snapshot Pool in which the LiveClone will reside. Ensure enough storage space is allocated to accommodate the LiveClone. If changes are needed, see the AGM online help for details on configuring Snapshot Pools.
11. Check **Mount for pre-processing** and the following screen provides the opportunity to mount the LiveClone to a host and then scrub the mounted image for sensitive data:



LiveClone Pre-Processing Options

Note: You may have to scroll down to see all options.

12. The **Mark data as non-sensitive** option only appears if the selected application is marked as containing sensitive data. Check this option if the application will not contain sensitive data after scrubbing. This will allow users who do not have access to sensitive data to access the Workflow and provision their work environments. See [Example User for Limited Workflow Access](#) on page 8 for details.
13. Select a host on which the LiveClone will be mounted.
14. Add a **Mount Drive** (Windows) or a **Mount Point** as needed if you want the image mounted to a location other than the default. If nothing is specified, then the Actifio Connector will assign a mount location.
15. Specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting the application. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
16. Specify a **Post Script** as needed. The post script in this example is used to scrub the application of sensitive information after it has been mounted. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
17. Specify timeouts long enough for the scripts to complete.
18. Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand and provision and re-provision any server to which they have access with the application data.
19. Click **Next** and the Mount Settings options are displayed. Use these options to limit what a user running this Workflow on demand can specify.



LiveClone Mount Settings

20. Specify a **Mount label** to make the mount easily identifiable.

21. Click the green **+** icon next to server names to specify the servers on which the application data will be automatically mounted. Users running the Workflow will be able to specify additional servers.
22. Specify a **Mount Mode, Mount Drive** as needed. Once specified, users running the Workflow will not be able to specify these mount settings.
23. **Pre Script** is a script used to prepare the environment prior to mounting or unmounting the final, scrubbed virtual application. Do not confuse this script with the script you specified to scrub data. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
24. Specify a script timeout long enough for the script(s) to complete.
25. **Post Script** is a script used to act on the virtual application after it has been scrubbed and mounted or unmounted. Do not confuse this script with the script you specified to scrub data. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
26. Specify a script timeout long enough for the script(s) to complete.

Note: *Scripts are usually used along with the **Remove mounted image after processing** option.*

27. **Do not** check Perform application aware mount.
28. **Remove mounted image after processing.** This is typically checked when you employ a script to process mounted data. Once the script finishes its task this option will unmount and delete the virtual application. See [Reporting](#) on page 2, [Data Warehousing](#) on page 2, and [Copying to Removable Media](#) on page 3.
29. Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand and provision and re-provision the servers to which they have access as specified in this section.

Direct Mount Workflow: Virtual Application

Direct Mount Workflows are used when the application image involved does not contain sensitive data. This allows the Workflow to mount an application's data as a virtual application or as just application data without first having to define a LiveClone and then scrub the LiveClone. Direct Mounts make application data and virtual applications available almost instantly.

In the following example, a Workflow is defined that:

- Generates new, or updates a virtual application from a single selected production image.
- Defines a schedule for updating the virtual application with the latest production data. Scheduled Workflows can also be run on demand.
- Mounts the data as a virtual application. A virtual application can be used and captured like any other application.
- Captures the mounted virtual application as a new application.

Users with the proper access rights can run this Workflow as described in [Direct Mount Run Now: Virtual Application](#) on page 36.

To define this Workflow:


1. Open the Actifio Application Manager.
2. In the Navigation pane, select the application to be mounted as a virtual application.
3. Click the **Workflow** tab and the Workflow Settings page is displayed:

If the selected application already has a Workflow, click **Add** and the Workflow Settings page is displayed:



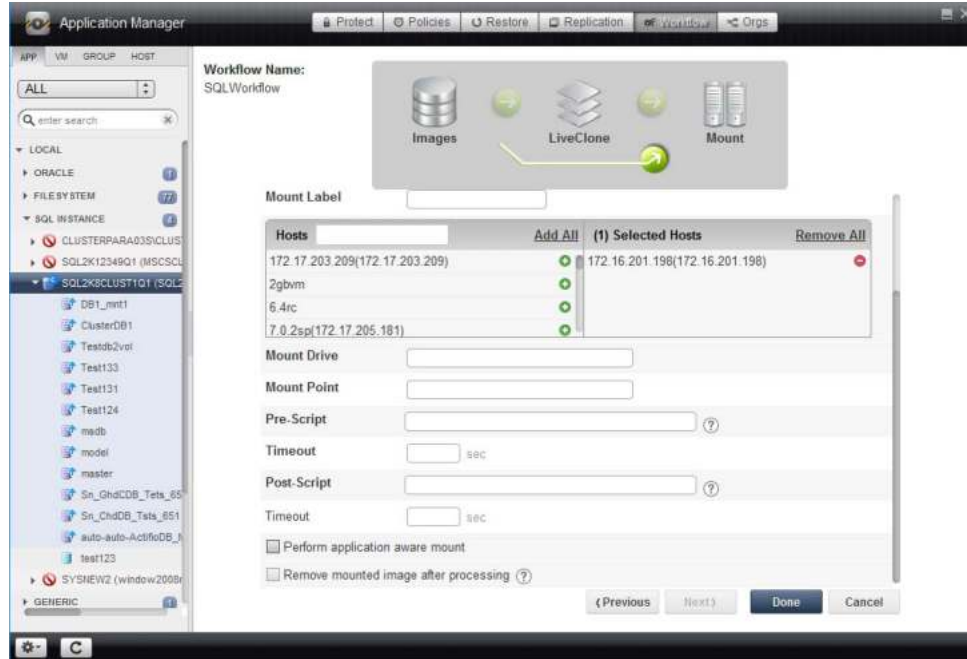
New Workflow Settings Page

4. Specify a name for the Workflow. Workflow names cannot include special characters.
5. Click **Direct Mount** and the Mount Settings page is displayed:



Mount Settings Options

6. For Schedule Type, select **Scheduled**. Scheduled Workflows can also be run on demand.
7. In **Schedule**, define a schedule to run the Workflow. The time selector uses a 24-hour clock.
8. For **Source Image**, select the latest snapshot image from the policy that produces the image to scrub.
9. Give the mount a label to make it easily identifiable.
10. In the **Hosts** section, click the green + icon to select the servers on which the data will be mounted. This ensures data will be mounted to those servers when the Workflow runs. The user who runs the Workflow will be given the option of mounting the data to other servers as needed.
11. Add a **Mount Drive** as needed if you want the image mounted to a location other than the default. If you do not specify a mount option, then the Actifio Connector will choose a mount location.

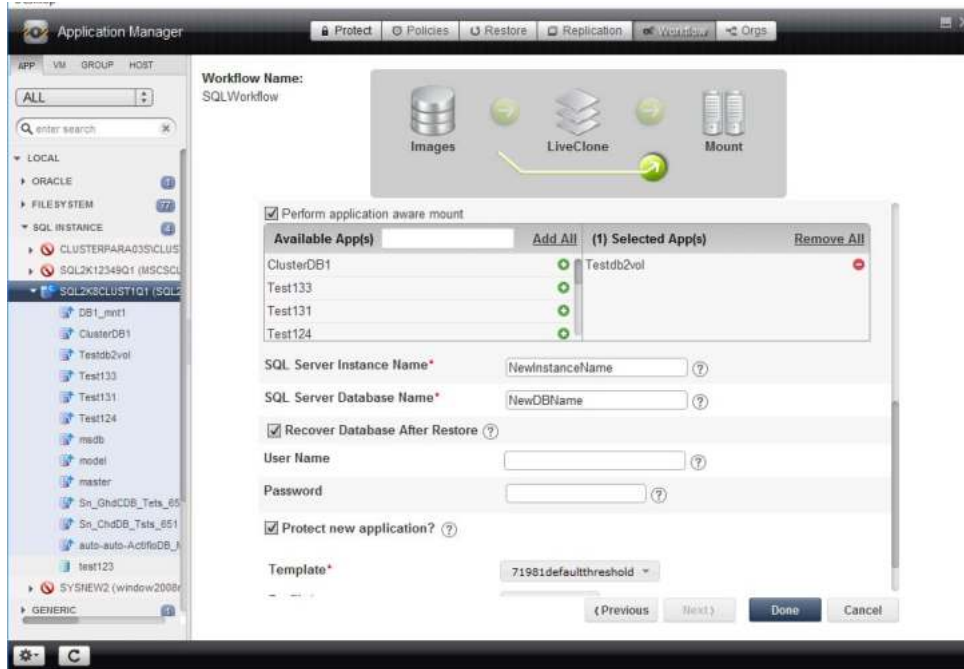


Mount Settings Options (Continued)

Note: You may have to scroll down to see all options.

12. Specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting or unmounting the application. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
13. Specify timeouts long enough for the script to complete.
14. Specify a **Post Script** as needed. The post script used to perform an operation on the data after it is mounted or unmounted, for example, initiating a copy to tape operation for long-term storage. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
15. Specify timeouts long enough for the script to complete.
16. Check **Perform application aware mount** and the application specific options to perform an application aware mount are displayed.

Note: The **Remove mounted image after processing** option disappears when you select **Perform application aware mount**. The **Remove mounted image after processing** option is typically used with pre and/or post scripts. It unmounts the image from target server, and deletes it after the script(s) finish. This option is typically used when scripts perform operations on the mounted data. See [Reporting](#) on page 2, [Data Warehousing](#) on page 2, and [Copying to Removable Media](#) on page 3.



Application Aware Mount Options

Note: You may have to scroll down to see the application aware mount options.

17. Click the green + next to the server name that will host the SQL Server Instance
18. In the space provided, enter a name for the virtual SQL Server Database.
19. 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.
20. In the spaces provided, enter a user name and password. The specified user must have the SQL Server Database role(s) necessary to perform all of the operations specified in the Workflow. For details on Microsoft SQL Server roles, see the AGM online help.
21. At this point you can:
 - Check the **Protect new application** option. This allows you to apply an SLA Policy Template and a Resource Profile to protect the virtual application as a new application. For details on SLA Policy Templates and Resource Profiles see the AGM online help.
 - Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand to provision and re-provision their environments with the virtual application.

Direct Mount Workflow: Application Data

Direct Mount Workflows are used when the application image involved does not contain sensitive data. This allows the Workflow to mount an application's data as a virtual application or as just application data without first having to define a LiveClone and then scrub the LiveClone. Direct Mounts make application data and virtual applications available almost instantly.

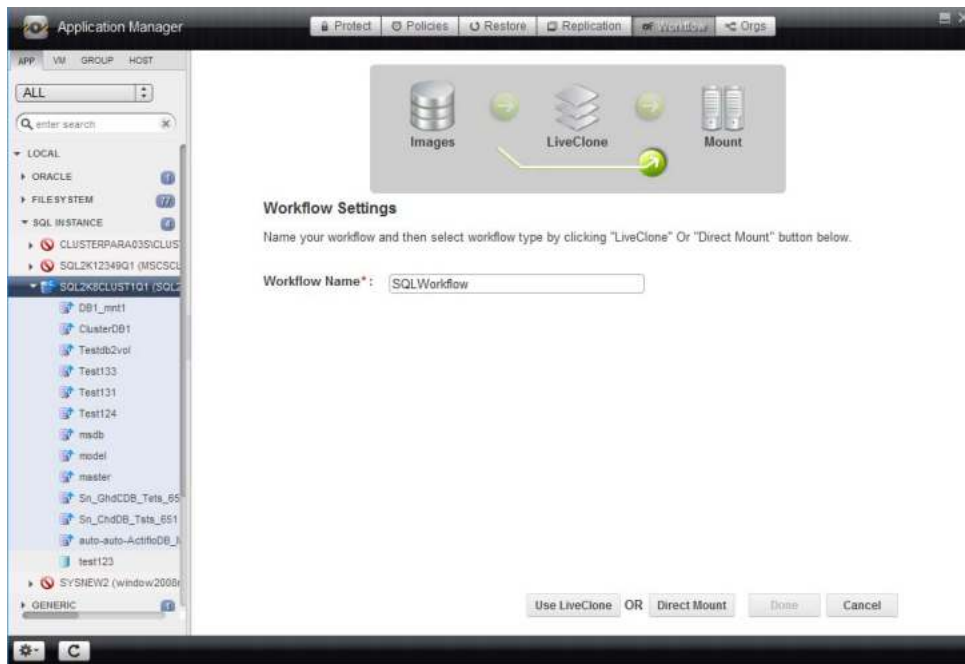
In the following example, a Workflow is defined that:

- Generates new, or updates mounted application data from a single selected production image.
- Defines a schedule for updating the application data with the latest production data. Scheduled Workflows can also be run on demand.
- Mounts just the application data.

Users with the proper access rights can run this Workflow as described in [Direct Mount Run Now: Application Data](#) on page 39.

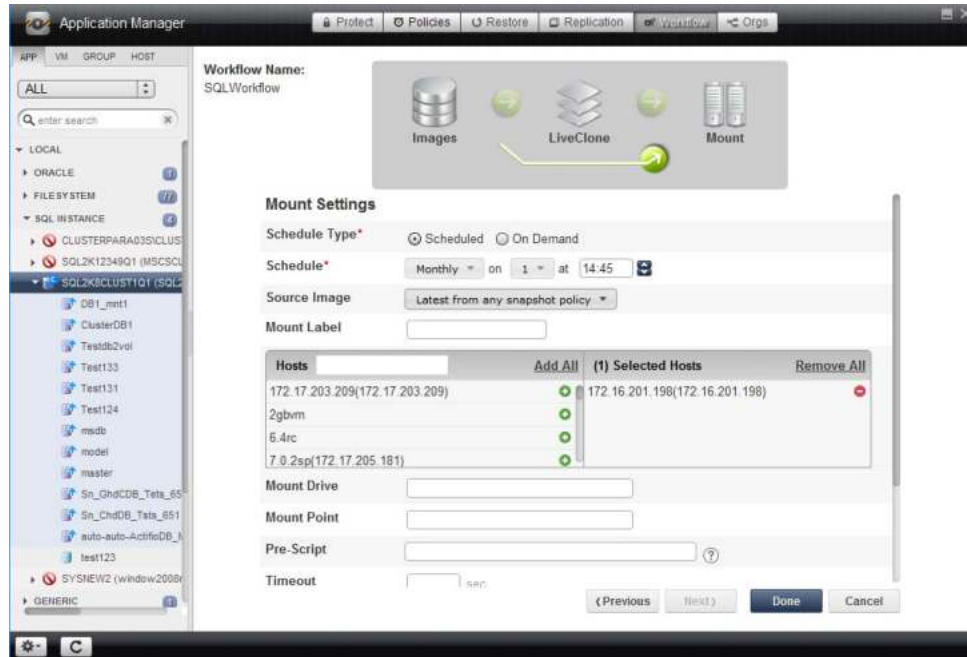
To define this Workflow:

1. Open the Actifio Application Manager.
2. In the Navigation pane, select the application.
3. Click the **Workflow** tab and the Workflow Settings page is displayed. If the application already has a Workflow, click **Add** and the Workflow Settings page is displayed:



New Workflow Settings Page

4. In the Workflow Settings page, specify a name for the Workflow. Workflow names cannot include special characters.
5. Click **Direct Mount** and the Mount Settings page is displayed:



Mount Settings Options

6. For Schedule Type, select **Scheduled**. Scheduled Workflows can also be run on demand.
7. In **Schedule**, define a schedule to run the Workflow. The time selector uses a 24-hour clock.
8. For **Source Image**, select the latest snapshot image from the policy that produces the image to scrub.
9. Give the mount a label to make it easily identifiable.
10. In the **Hosts** section, click the green + icon to select the servers on which the data will be mounted. This ensures data will be mounted to those servers when the Workflow runs. The user who runs the Workflow will be given the option of mounting the data to other servers as needed.

Note: For scheduled Direct Mount Workflows that just mount the data, you must select at least one host.

11. Add a **Mount Drive** as needed if you want the image mounted to a location other than the default. If you do not specify a mount, then the Actifio Connector will choose a mount location.

Note: You may have to scroll down to see all options.

12. Specify a **Pre Script** as needed. The pre script is used to configure the environment prior to mounting the application. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
13. Specify a **Post Script** as needed. The post script is used to perform an operation on the data after it is mounted, for example, initiating a copy to tape operation for long-term storage. This script must reside in a folder named /act/scripts on the server that will host the mounted image. See [Chapter 5, Workflow Pre and Post Scripts](#) for scripting details.
14. If you specified scripts, specify timeouts long enough for the scripts to complete.
15. Because you are only mounting the data, you will not select **Perform application aware mount**.

16. **Remove mounted image after processing.** This is typically checked when you employ a script to process mounted data. Once the script finishes its task, this option will unmount and delete the virtual application. See [Reporting](#) on page 2, [Data Warehousing](#) on page 2, and [Copying to Removable Media](#) on page 3.
17. Click **Done** and the Workflow will run as scheduled. Users with proper access can run this Workflow on demand to provision and re-provision their environments with the application data.

4 Running Workflows On Demand

Workflows can be configured to run according to a schedule or run on-demand only. Scheduled Workflows can be run on demand.

Running a Workflow on demand is particularly useful when the Workflow's data source has been updated but the Workflow is not scheduled to run for a while. The Workflow can be run on demand to provision/re-provision an environment without having to wait for the next scheduled run of the Workflow.

To run Workflows on demand:

- Users must be assigned a role with the necessary rights, and assigned to the proper organizations to access an Actifio appliance. For details on Roles, Organizations, and Users, see [Configuring Roles, Organizations, and Users](#) on page 5.
- User's rights also determine whether they can update a LiveClone as part of running a LiveClone Workflow.
- The options that were or were not specified when the Workflow was defined determines which options a user can specify when they run a Workflow.

This chapter provides step-by-step examples for:

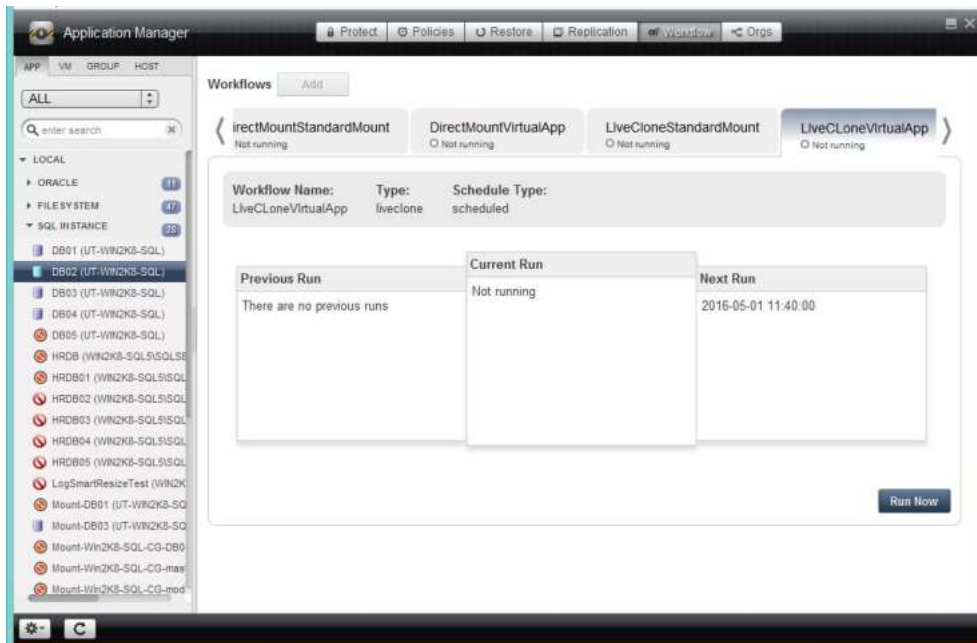
- [LiveClone Run Now: New Scrubbed Virtual Application](#) on page 28
- [LiveClone Run Now: Re-Provisioning a Scrubbed Virtual Application](#) on page 31
- [LiveClone Run Now: Scrubbed Application Data](#) on page 34
- [Direct Mount Run Now: Virtual Application](#) on page 36
- [Direct Mount Run Now: Application Data](#) on page 39

LiveClone Run Now: New Scrubbed Virtual Application

This section shows how to run a LiveClone Workflow like the one created in [LiveClone Workflow: Scrubbed Virtual Application](#) on page 10.

Running this Workflow allows a user with the proper privileges to provision a selected server with a new virtual application.

1. Open the Application Manager to the **Workflow** tab.
2. In the Navigation pane, select the application that will be used to provision the virtual application. The associated Workflow schedule appears in the Workflow screen.




Workflow For Selected Application

3. Click **Run Now** and the Run Now Settings page is displayed.
4. Users assigned **Manage LiveClone** rights in their assigned role will have the option of selecting a new image for the Workflow's LiveClone. Such users can check the **Refresh LiveClone** option and from the drop down menu, select an image to be scrubbed.

Note: Users assigned a role with **Manage LiveClone** rights must also have access to sensitive data. See [Example User for Limited Workflow Access](#) on page 8.

At this point, if you click **Done**, then only the LiveClone will be updated.



Run Now Settings

5. Check **Create/Refresh a mounted image or virtual application** and more Run Now Settings are displayed:



Run Now Settings

6. From the drop down menu, select the source (scrubbed) image to use for the virtual application.
7. If the application is a database and it is captured along with its logs, then the log recovery pane appears with a slider above the target host selection pane. Select whether to use the time of the host or the time of your workstation, and then select the recovery point by using the slider, the calendar and arrows, or by simply typing a value into the field.
8. If the creator of the Workflow did not specify the host to be provisioned in the Workflow, click the green + icon next to a host name to add it to the list of hosts to be provisioned. When provisioning/re-provisioning a virtual application, only one host can be selected.



Application Aware Mount Options

Note: You may have to scroll down to see all options.

If the creator of the Workflow did not specify the options for the virtual application:

9. Check the box for **Perform application aware mount**.
10. Select **Create New Virtual Application**.
11. Select an SQL Server Instance on the selected host.
12. In the space provided, enter a name for the new virtual SQL Server Database.
13. The username and password is only required if the Actifio Connector used does not have privileges to apply logs or detach the database during an unmount operation.
14. Click **Done** and the new virtual application will be mounted as defined. Access to the virtual application is controlled by the server on which it is mounted.

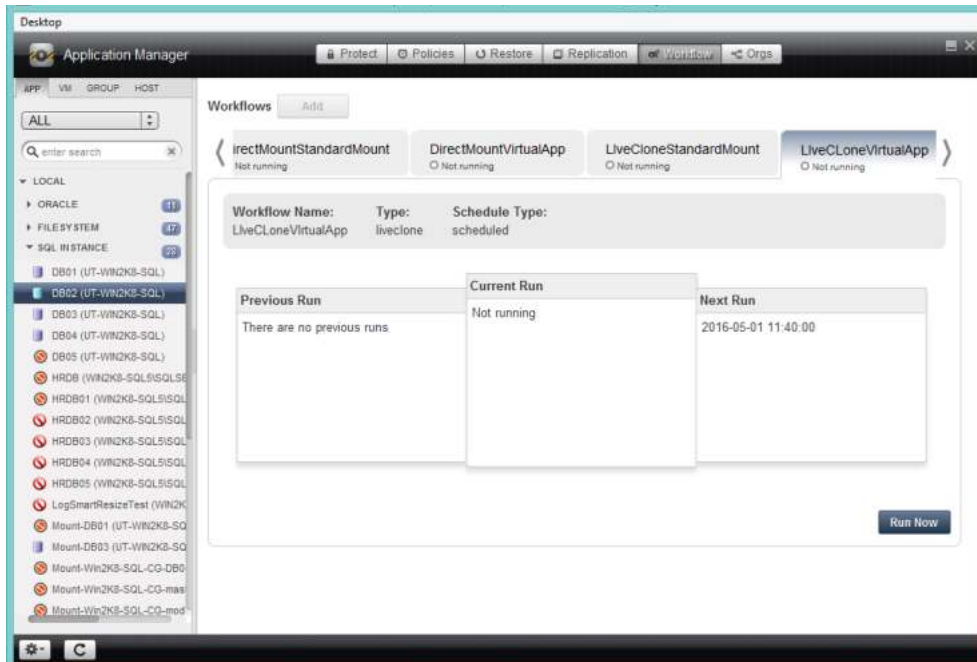
LiveClone Run Now: Re-Provisioning a Scrubbed Virtual Application

This section shows how to run a LiveClone Workflow like the one created in [LiveClone Workflow: Scrubbed Virtual Application](#) on page 10.

Running this Workflow allows a user with the proper privileges to re-provision a server with an updated virtual application.

To run a Workflow to refresh a previously provisioned virtual application with a new image:

1. Open the Application Manager to the **Workflow** tab.
2. In the Navigation pane, select the application. The Workflows available for the application and their run schedules appear in the Workflow screen.



Workflow For Selected Application

3. Select the Workflow to execute and click **Run Now**.
4. Users assigned **Manage LiveClone** rights in their assigned role will have the option of selecting a new image for the Workflow's LiveClone. Such users can check the **Refresh LiveClone** option and from the drop down menu, select an image to be scrubbed.

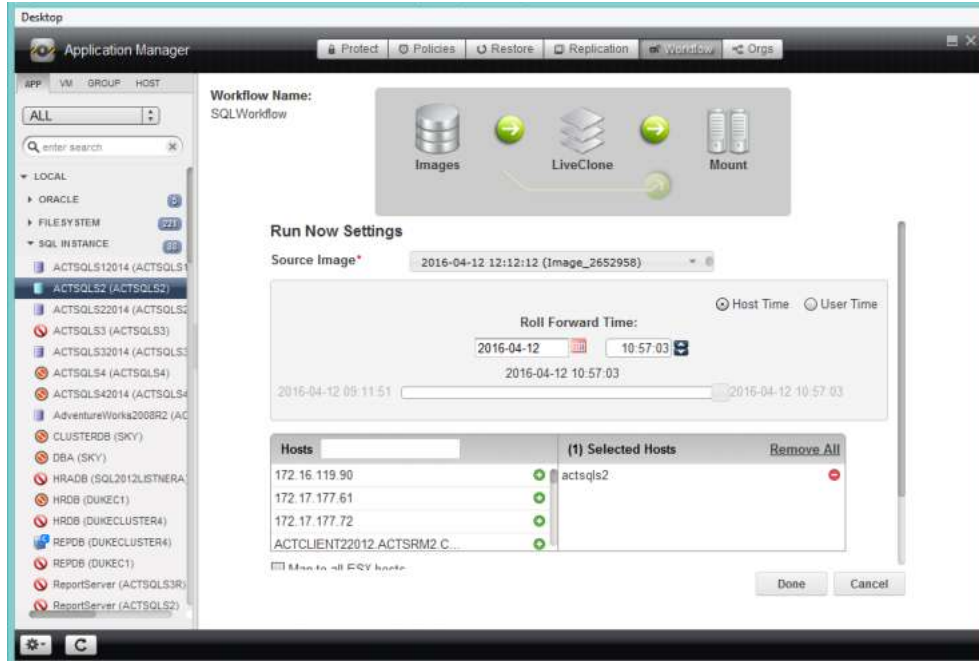
Note: Users assigned a role with **Manage LiveClone** rights must also have access to sensitive data. See [Example User for Limited Workflow Access](#) on page 8.

At this point, if you click **Done**, then only the LiveClone will be updated.



Run Now Settings

5. Check **Create/Refresh a mounted image or virtual application** and more Run Now Settings are displayed:



Run Now Settings

6. In the Run Now Settings screen, select the **Source (scrubbed) Image** to use for provisioning.
7. If the application is a database and it is captured along with its logs, then the log recovery pane appears with a slider above the target host selection pane. Select whether to use the time of the host or the time of your workstation, and then select the recovery point by using the slider, the calendar and arrows, or by simply typing a value into the field.
8. If the creator of the Workflow did not specify the host to be provisioned in the Workflow, click the green + icon next to a host name to add it to the list of hosts to be provisioned. When provisioning/re-provisioning a virtual application, only one host can be selected.



Run Now Settings (Continued)

Note: You may have to scroll down to see all options.

9. Check the box for **Perform application aware mount**.
10. Select **Refresh Existing Virtual Application**.
11. At the **Application** drop-down, select the application to refresh.
12. Click **Done** and the new virtual application will be mounted as defined by the Workflow. Access to the virtual application is controlled by the server on which it is mounted.

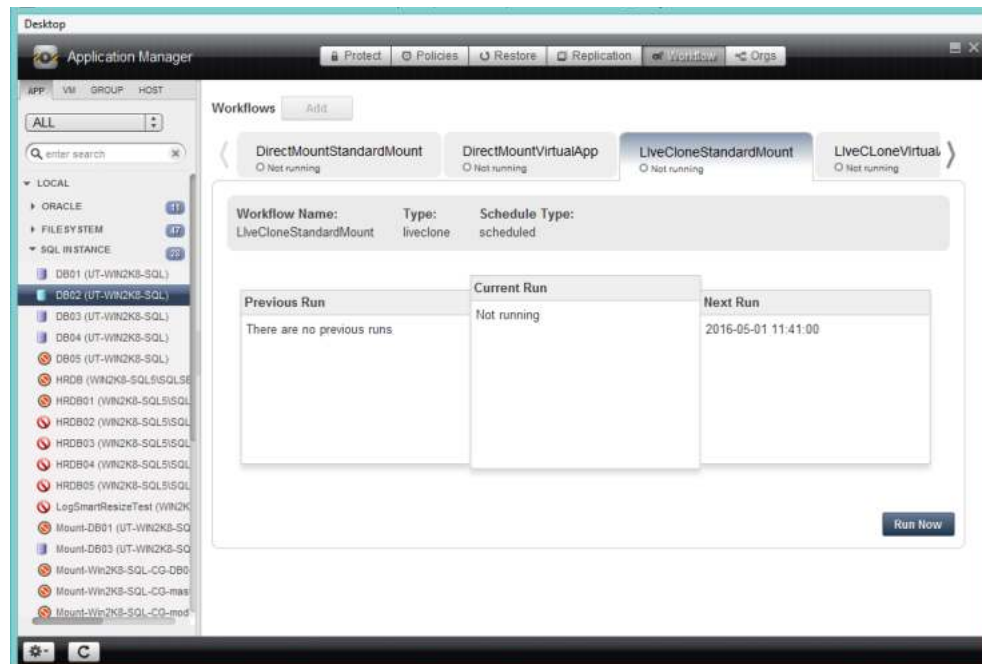
LiveClone Run Now: Scrubbed Application Data

This section shows how to run a LiveClone Workflow like the one created in [LiveClone Workflow: Scrubbed Application Data](#) on page 16.

Running this Workflow allows a user with the proper privileges to provision servers with scrubbed data from an application.

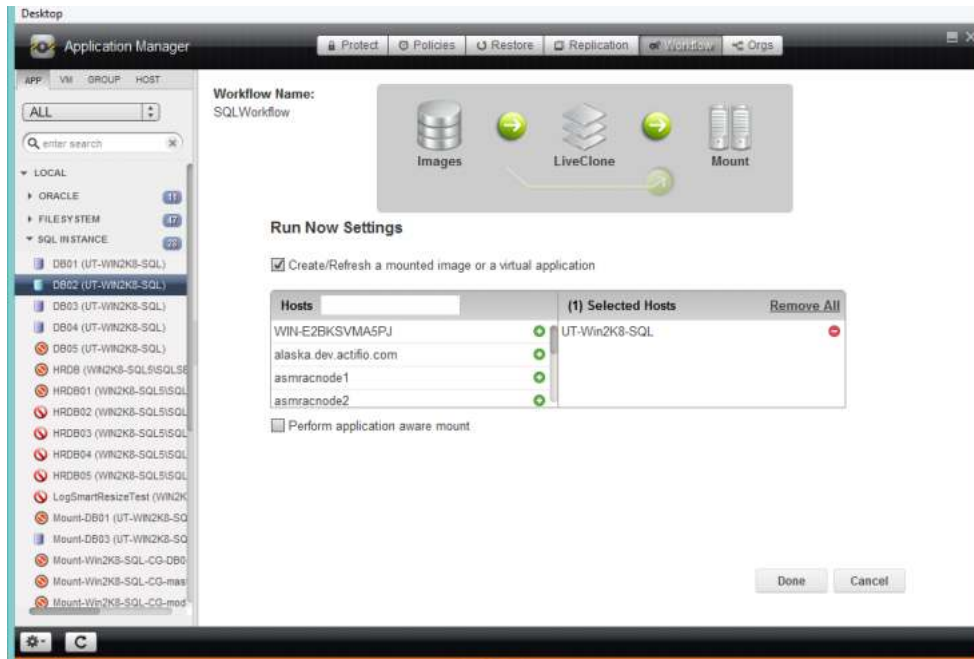
To provision an environment with scrubbed data from an application:

1. Open the Application Manager to the **Workflow** tab.
2. In the Navigation pane, select the application used to provision the environment. The associated Workflow schedule appears in the Workflow screen.



Workflow For Selected Application

3. Click **Run Now** and the Run Now Settings page is displayed:



Run Now Settings

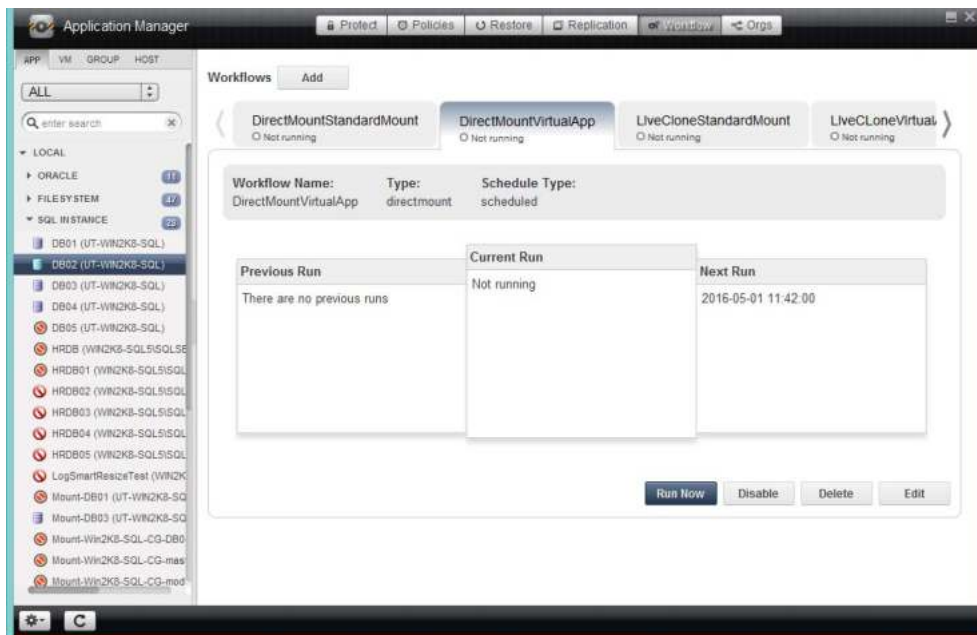
4. Select the **Source Image** to use for provisioning.
5. If the application is a database and it is captured along with its logs, then the log recovery pane appears with a slider above the target host selection pane. Select whether to use the time of the host or the time of your workstation, and then select the recovery point by using the slider, the calendar and arrows, or by simply typing a value into the field.
6. If the creator of the Workflow did not specify all the hosts to be provisioned in the Workflow, click the green **+** icon next to a host name to add it to the list of hosts to be provisioned.
7. Click **Done** and the application data will be mounted as defined. Access to the data is controlled by the server on which it is mounted.

Direct Mount Run Now: Virtual Application

This section shows how to run a Direct Mount Workflow like the one created in [Direct Mount Workflow: Virtual Application](#) on page 20.


Running this Workflow allows a user with the proper privileges to provision a selected server with a new virtual application.

1. Open the Application Manager to the **Workflow** tab.
2. In the Navigation pane, select the application that will be used to provision the new virtual application. The associated Workflow schedule appears in the Workflow screen.



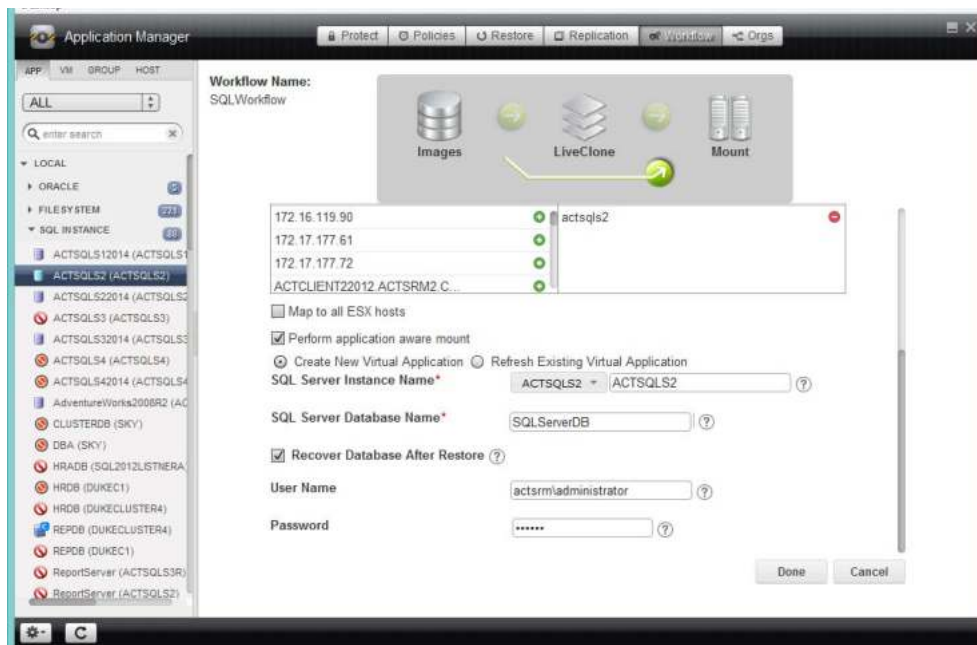
Workflow For Selected Application

3. Click **Run Now** and the Run Now Settings page is displayed:



Run Now Settings

4. From the drop down menu, select the source image to use for the virtual application.
5. If the application is a database and it is captured along with its logs, then the log recovery pane appears with a slider above the target host selection pane. Select whether to use the time of the host or the time of your workstation, and then select the recovery point by using the slider, the calendar and arrows, or by simply typing a value into the field.
6. If the creator of the Workflow did not specify the host to be provisioned in the Workflow, click the green + icon next to a host name to add it to the list of hosts to be provisioned. When provisioning/re-provisioning a virtual application, only one host can be selected.



Run Now Settings (Continued)

Note: You may have to scroll down to see all options.

7. Check the box for **Perform application aware mount**.

If you are re-provisioning an environment:

- o Select **Refresh Existing Virtual Application**.
- o Click **Done** and the virtual application will be updated with the new data.

If you are provisioning a new environment:

- o Select **Create New Virtual Application**.

If the creator of the Workflow did not specify the options for the virtual application:

- o Select an SQL Server Instance on the selected host.
- o In the space provided, enter a name for the new virtual SQL Server Database. The username and password is only required if the Actifio Connector used does not have privileges to apply logs or detach the database during an unmount operation.

8. Click **Done** and the new virtual application will be mounted as defined. Access to the virtual application is controlled by the server on which it is mounted.

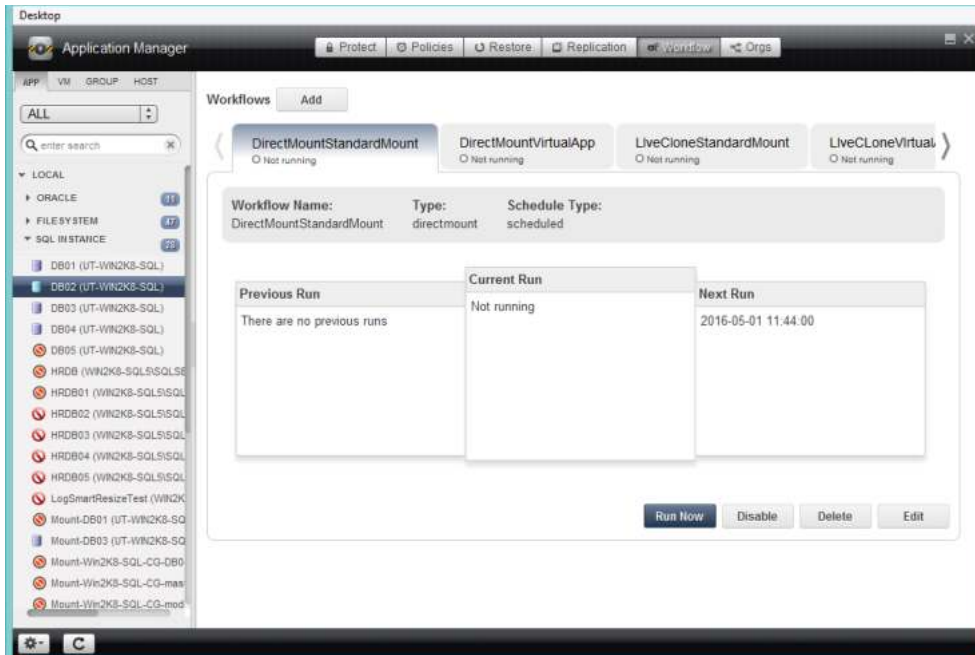
Direct Mount Run Now: Application Data

This section shows how to run a Direct Mount Workflow like the one created in [Direct Mount Workflow: Application Data](#) on page 24.

Running this Workflow allows a user with the proper privileges to provision servers with just the data from an application.


To provisioned an environment with just the data from an application:

1. Open the Application Manager to the **Workflow** tab.
2. In the Navigation pane, select the application used to provision the environment. The associated Workflow schedule appears in the Workflow screen.



Workflow For Selected Application

3. Click **Run Now** and the Run Now Settings page is displayed:



Run Now Settings

4. Select the **Source Image** to use for provisioning.
5. If the application is a database that is captured along with its logs, then the log recovery pane appears with a slider above the target host selection pane. Select whether to use the time of the host or the time of your workstation, and then select the recovery point by using the slider, the calendar and arrows, or by simply typing a value into the field.
6. If the creator of the Workflow did not specify all of the hosts to be provisioned in the Workflow, click the green + icon next to a host name to add it to the list of hosts to be provisioned. When provisioning/re-provisioning application data, only one host can be selected.
7. Click **Done** and the application data will be mounted as defined. Access to the data is controlled by the server on which it is mounted.

5 Workflow Pre and Post Scripts

This chapter includes:

- [Introduction](#) on page 41
- [Environment Variables](#) on page 42
- [High-Level Examples](#) on page 43
- [Variables Used to Mount Databases](#) on page 44
- [In All Cases](#) on page 48
- [Example Script](#) on page 49

Introduction

Workflows mount and unmount captured images according to a schedule or on demand. In a Workflow you can call:

- A pre script that runs before an image is mounted and/or unmounted
- A post script that runs after an image is mounted and/or unmounted

Scripts before and after data is mounted or unmounted allow you to perform operations on the data such as:

- Scrubbing for sensitive information
- Generating reports
- Warehousing, especially when extracting, transforming, and loading (ETL) is required
- Copying to removable media

The pre and post scripts must reside in a folder on the server that hosts the mounted Workflow image:

- Windows: C:\Program Files\Actifio\Scripts\
- Linux: /act/scripts

Note: For Workflows that involve a database with its logs, configure pre and post scripts that are not applied to the logs. See [Example Script](#) on page 49 for a script example that contains a command that will skip a database's log.

When using PRE or POST scripts to provision a new virtual database, it is important to include logic to detect what operation triggered the script's execution.

This logic is necessary because pre and post scripts are run not only during the original database provisioning and during subsequent refreshes, but also during restores of snapshots of the virtual database (i.e. rolling the virtual database back to a previous point in time).

It is common for these scripts to perform some database preparation that would not be needed or applicable during a restore from a snapshot of an earlier point in time.

To determine if a script was run as a result of a restore from a snapshot of the virtual database, check for the environment variable `ACT_MOUNT_REPLACERESTORE` to be set to a value of `true` and avoid running any processes that are not applicable in this situation.

Environment Variables

Environment variables are set by the Actifio Connector before launching pre/post scripts. These environment variables provide information about the tasks being performed by the Connector. When you create pre/post scripts, you can employ these environment variables to implement the script's desired logic. Actifio Connector-defined environment variables are prefixed with `ACT_`.

For example, an environment variable `ACT_APPNAME` indicates that the Connector is performing some operation for an application pointed to by the environment variable.

These are common environment variables, with sample values in parentheses.

- `JOBNAME`: The name of the job (e.g. `Job_0123456`)
- `APPID`: The ID of the application (e.g. `4186`)
- `APPNAME`: Name of the application (e.g. `My-DB`)
- `HOSTNAME`: The name of the host which is the target of this job (e.g. `Jupiter`)
- `SOURCEHOST`: The name of the host that was the source for this application (e.g. `Saturn`)
- `JOBTYPE`: a text version of the job class (e.g. `mount, unmount`)
- `PHASE`: A text string that describes the job phase (e.g. `pre, post`)
- `TIMEOUT`: Define the duration of the script, how long the script is allowed to run
- `OPTIONS`: Policy options that apply to this job
- `ACT_MOUNT_REPLACERESTORE`: If set to `true`, indicates script was run as the result of a restore operation.

High-Level Examples

The following examples are high-level descriptions of some of the more common pre and post script combinations.

Provision a New Database, Log Capture Not Enabled

This example is the most common use. There is a single action: Mount the database image. In this example, you will:

- Call Pre-Script before performing the database image mount
- Call the Post-Script after performing the database image mount and performing any app-aware activities (i.e. call the script after the database is online when using an app-aware mount).

Refresh a Previously Provisioned Database, Log Capture Not Enabled

Instead of just performing a mount, the script performs an unmount of the previous database image, then a mount of the new one. In this example, you will:

- Call Pre-Script before unmounting the existing database
- Call Post-Script after unmounting the existing database
- Call Pre-Script before mounting the new database
- Call Post-Script after mounting the new database and doing app-aware activities

Most often if there is a problem with this script, it is usually an issue with the script call during Post Unmount. This is usually because the Post-Script is doing some database work and fails when run as a Post Unmount and there is no database to work on at that time.

Provision a New Database with Log Capture Enabled

Capturing the database and its logs add another level of complexity to the scripting effort. In this example, you will:

- Call the Pre-Script before mounting the database logs
- Call the Post-Script after mounting the database logs
- Call the Pre-Script before mounting the database image
- Call the Post-Script after mounting the database image and performing app-aware activities i.e. after the database is online and has been rolled forward to the desired point-in-time using the logs.
- Call the Pre-Script before unmounting the database logs
- Call the Post-Script after unmounting the database logs

In this very common scenario, the Pre-Script and the Post-Script are run three times. If you run into a problem it will most likely be the first call to the Post-Script. The first Post-Script performs some database action, but the first time it is run the database is not yet on line and the logs have been mounted.

Variables Used to Mount Databases

When a Pre or Post script is called as a result of a mount or unmount, several environment variables are set before the scripts is actually executed. This is done to help the script author understand the circumstances under which the script was called.

Variables Set for a Database Mount Without Logs

- ACT_APPNAME=bigdb
- ACT_JOBNAME=Job_37204210
- **ACT_JOBTYPE=mount**
- ACT_MULTI_END=true
- **ACT_MULTI_OPNAME=mount**
- **ACT_PHASE=pre**
- ACT_POLICY=Production to Snap 1
- ACT_PROFILE=AWS Profile
- ACT_SCRIPT_TMOUT=600
- ACT_SOURCEHOST=hq-sql
- ACT_TEMPLATE=StreamSnap
- PHASE=pre
- dbname=bigdbtest
- recover=true
- sqlinstance=DEMO-SQL-11

The highlighted variables ACT_JOBTYPE, ACT_MULTI_OPNAME, and ACT_PHASE show this is a run of a Pre-Script on a mount operation on a mount job.

Pay particular attention to the difference between ACT_JOBTYPE vs. ACT_MULTI_OPNAME. The value for ACT_JOBTYPE may be "mount", "unmount", or "reprovision". This is the job type shown in the job monitor.

However, for a reprovision job, there will be an unmount and a mount job. The variable ACT_MULTI_OPNAME was created to tell the script which specific action is being taken when a job will perform multiple actions, like the case with reprovision jobs. ACT_MULTI_OPNAME will always be "mount" or "unmount" and is typically the better choice to use when checking what actions a script should perform.

Variables Set for a Database Reprovision With Logs

1. Unmount db pre
 - o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=db**
 - o **ACT_MULTI_END=true**
 - o **ACT_MULTI_OPNAME=unmount**
 - o **ACT_PHASE=pre**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600

- o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=pre
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11
2. Unmount db post
- o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=db**
 - o ACT_MOUNT_POINTS=C:\actifio\smalldb\;
 - o **ACT_MULTI_END=true**
 - o **ACT_MULTI_OPNAME=unmount**
 - o **ACT_PHASE=post**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=post
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11
3. Mount log pre
- o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=log**
 - o ACT_MOUNT_REPROVISION=true
 - o **ACT_MULTI_OPNAME=mount**
 - o **ACT_PHASE=pre**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=pre
 - o dbname=smalldbtest

- o recover=true
 - o sqlinstance=DEMO-SQL-11
4. Mount log post
- o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=log**
 - o ACT_MOUNT_POINTS=C:\Windows\act\Job_37201405_mountpoint_150169372132911\;
 - o ACT_MOUNT_REPROVISION=true
 - o **ACT_MULTI_OPNAME=mount**
 - o **ACT_PHASE=post**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=post
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11
5. Mount db pre
- o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=db**
 - o ACT_MOUNT_REPROVISION=true
 - o **ACT_MULTI_END=true**
 - o **ACT_MULTI_OPNAME=mount**
 - o **ACT_PHASE=pre**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=pre
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11

6. Mount db post
 - o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=db**
 - o ACT_MOUNT_POINTS=C:\actifio\smalldb\;
 - o ACT_MOUNT_REPROVISION=true
 - o **ACT_MULTI_END=true**
 - o **ACT_MULTI_OPNAME=mount**
 - o **ACT_PHASE=post**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=post
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11
7. Unmount log pre
 - o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**
 - o **ACT_LOGSMART_TYPE=log**
 - o **ACT_MULTI_OPNAME=unmount**
 - o **ACT_PHASE=pre**
 - o ACT_POLICY=12hr Snap
 - o ACT_PROFILE=Remote Profile HQ to DR
 - o ACT_SCRIPT_TMOUT=600
 - o ACT_SOURCEHOST=hq-sql
 - o ACT_TEMPLATE=Silver LogSmart
 - o PHASE=pre
 - o dbname=smalldbtest
 - o recover=true
 - o sqlinstance=DEMO-SQL-11
8. Unmount log post
 - o ACT_APPNAME=smalldb
 - o ACT_JOBNAME=Job_37201405
 - o **ACT_JOBTYPE=reprovision**

- o **ACT_LOGSMART_TYPE=log**
- o ACT_MOUNT_POINTS=C:\Windows\act\Job_37201405_mountpoint_150169372132911\;
- o **ACT_MULTI_OPNAME=unmount**
- o **ACT_PHASE=post**
- o ACT_POLICY=12hr Snap
- o ACT_PROFILE=Remote Profile HQ to DR
- o ACT_SCRIPT_TMOUT=600
- o ACT_SOURCEHOST=hq-sql
- o ACT_TEMPLATE=Silver LogSmart
- o PHASE=post
- o dbname=smalldbtest
- o recover=true
- o sqlinstance=DEMO-SQL-11

Variables of interest are in bold and include:

- o ACT_JOBTYPE (mount, unmount, reprovision)
- o ACT_MULTI_OPNAME (mount, unmount)
- o ACT_PHASE (pre, post)
- o ACT_MULTI_END (true) - not always present
- o ACT_LOGSMART_TYPE (db, log) - not always present

In All Cases

Check

- ACT_MULTI_END. Only take action when this variable exists, and is set to true. Otherwise you are taking an action during a log volume mount or unmount, and will almost certainly end up running your logic at the wrong time.
- ACT_MULTI_OPNAME. This will always be unmount or mount. The two most common times you will want to run scripts are on POST scripts when ACT_MULTI_OPNAME is set to mount, or on PRE scripts when ACT_MULTI_OPNAME is set to unmount (i.e. before removing the database during an unmount portion of a refresh).

You very rarely if ever will need to check the following:

- You do not need to check ACT_LOGSMART_TYPE if you check ACT_MULTI_END
- You do not need to check ACT_JOBTYPE if you check ACT_MULTI_OPNAME)

Example Script

The following script example uses three environment variables:

- **ACT_JOBTYPE** - Identifies whether the job is a mount or unmount operation
- **ACT_PHASE** - Identifies whether the phase is either pre or post
- **ACT_MULTI_END** - Used only if both a database and its log are mounted. When this is "true" the database is in a state where it can be accessed.

Note: The following example is not for use in a production environment. It is not warranted by Actifio.

```
@echo off
setlocal enabledelayedexpansion
echo "*** Running user script: Job - %ACT_JOBNAME% Type - %ACT_JOBTYPE% Phase - %ACT_PHASE% ***"
set desiredJobType=true
REM Use the first if clause to perform application specific operations during mount and in this
example, scrub-mount operation.
REM Use the second if clause to perform any application specific operation during unmount and in
this example, scrub-unmount operation.
REM if "%ACT_JOBTYPE%" neq "mount" if "%ACT_JOBTYPE%" neq "scrub-mount" set
desiredJobType=false
if "%ACT_JOBTYPE%" neq "unmount" if "%ACT_JOBTYPE%" neq "scrub-unmount" set
desiredJobType=false
if "%desiredJobType%" neq "true" (
echo "NO-OP for job type %ACT_JOBTYPE%"
goto :END
)
REM Use the first if clause to perform application specific operations during the pre phase
REM Use the second if clause to perform application specific operations during the post phase
REM if "%ACT_PHASE%" neq "post" (
if "%ACT_PHASE%" neq "pre" (
echo "NO-OP for phase %ACT_PHASE%"
goto :END
)
)
REM For multi-phase jobs (database and logs) check if the database has been mounted and the logs
applied then skip logs.
REM If the operation needs to be performed in phases other than the last phase, modify the
clause
if "%ACT_MULTI_END%" neq "" if "%ACT_MULTI_END%" neq "true" (
echo "NO-OP for multi-phase operation"
goto :END
)
cd "C:\Program Files\Actifio\scripts"
echo "*** Running application specific logic: Job - %ACT_JOBNAME% Type - %ACT_JOBTYPE% Phase -
%ACT_PHASE% ***"
REM Any application specific commands will go here
echo "*** Finished running application specific logic : Job - %ACT_JOBNAME% Type -
%ACT_JOBTYPE% Phase - %ACT_PHASE% ***"
:END
endlocal
```

