
Getting Started with Actifio Copy Data Management

Copyright, Trademarks, and other Legal Matter

Copyright © 2009 - 2019 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

Preface	v
The ActifioNOW Customer Portal	v
Actifio Support Centers	v
Chapter 1 - Introduction	1
Chapter 2 - Basic Concepts	3
Actifio Appliances.....	3
Local and Remote Appliances.....	3
Actifio CDS and CDX Physical Appliances.....	3
Actifio Sky Virtual Appliance.....	3
Actifio Sky for AWS Cloud Appliance.....	3
VDisks.....	4
VDisk Consumption	4
Storage Pools.....	5
Snapshot Pool	6
Out-of-Band Data and Staging Disks.....	7
In-Band Data.....	9
Dedup Pool	11
How Dedup Pool Space is Consumed and Reclaimed: Garbage Collection	12
Policy Templates and Policies	14
Resource Profiles.....	14
Data Capture Methods	15
Out-of-Band.....	15
In-Band.....	15
LAN-Free.....	16
Capture Mechanisms	17
Capture Options	18
Capture Individual and Groups of Applications	18
Capture Application Data in Actifio Consistency Groups	18
Capture a VM's Applications and Boot Volume	19
Capture Entire VMware and Hyper-V VMs.....	19
Capture Generic Applications	19
Capture Database Logs.....	20

Replicating Data	20
Dedup Backup to Dedup DR Replication.....	21
Sending Snapshots to an OnVault Pool's Defined Storage	21
Production to Mirror Policy Replication	21
Actifio NAS Director.....	23
Managed Data License.....	23
Chapter 3 - Capturing Applications Overview	25
Adding Hosts	26
Discovering Applications and VMs	27
Creating Policy Templates.....	28
Creating Policies.....	30
Production to Snapshot Policies.....	31
Snapshot to Dedup Backup Policies.....	31
Dedup Backup to Dedup DR Policies.....	32
Snapshot to OnVault Policies.....	32
Production to Mirror Policies	33
Production Direct to Dedup Policies.....	34
Resource Profiles.....	35
Applying Policy Templates and Resource Profiles	36
Chapter 4 - Accessing Data	37
Appendix A - Glossary of Terms	39

Preface

The information presented in this guide is intended for users who are new to managing and accessing data with an Actifio CDS or Sky appliance. This document assumes that the Actifio appliance(s) have been installed and are ready to begin managing your data.

Your Actifio appliance's Documentation Library contains detailed, step-by-step, application-specific instructions on how to protect and access your data. Each guide is in PDF format and may be viewed online, downloaded, or printed on demand. The following guides will be of particular interest:

- ***Setting Up Users and Roles With the Domain Manager***
- ***Configuring Resources and Settings With the Domain Manager***
- ***Connecting Hosts to Actifio Appliances***
- ***Planning and Developing Service Level Agreements***
- ***Virtualizing and Protecting Copy Data with the Application Manager***
- ***Accessing and Recovering Copy Data with the Application Manager***
- ***Replicating Data Using Actifio Appliances***

The ActifioNOW Customer Portal

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

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

To log into the ActifioNOW customer portal:

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

Actifio Support Centers

To contact an Actifio support representative, you can:

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

From anywhere: +1.315.261.7501

US Toll-Free: +1.855.392.6810

Australia: 0011 800-16165656

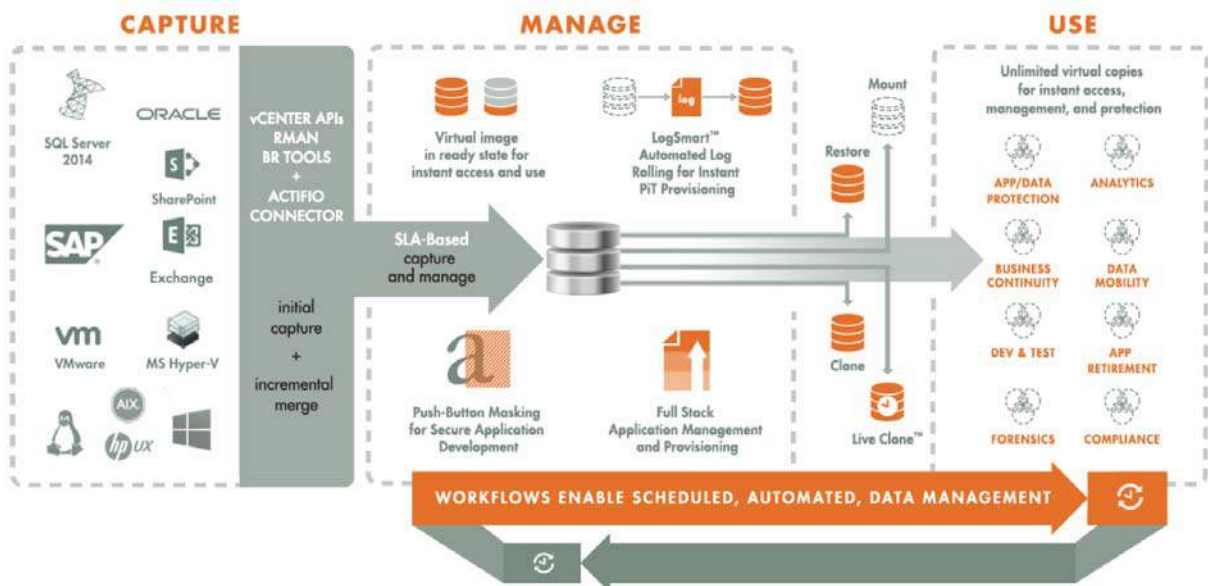
Germany: 00 800-16165656

New Zealand: 00 800-16165656

UK: 0 800-0155019

1 Introduction

Actifio CDS and Sky appliances are highly scalable copy data management platforms that virtualize application data to improve the resiliency, agility, and cloud mobility of your business. They virtualize data in much the same way other technologies have virtualized servers and networks. Actifio appliances enable 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.



Actifio Copy Data Management

Application data is captured at the block level, in native format, according to a specified SLA. A golden copy of that data is created, moved, 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 be accessed instantly for use, without proliferating physical copies and taking up additional storage infrastructure.

2 Basic Concepts

This chapter introduces the concepts of how an Actifio appliance captures, manages, and accesses data. Understanding these concepts will help you to be successful with Actifio copy data management.

- [Actifio Appliances](#) on page 3
- [VDisks](#) on page 4
- [Storage Pools](#) on page 5
- [Policy Templates and Policies](#) on page 14
- [Resource Profiles](#) on page 14
- [Data Capture Methods](#) on page 15
- [Capture Mechanisms](#) on page 17
- [Capture Options](#) on page 18
- [Replicating Data](#) on page 20
- [Actifio NAS Director](#) on page 23
- [Managed Data License](#) on page 23

Actifio Appliances

Actifio appliances capture and manage data locally and can replicate protected data to other Actifio appliances.

Local and Remote Appliances

Multiple Actifio appliances can be joined in any combination of primary and secondary relationships. An exchange of certificates is required to join appliances. Once joined, application data can be replicated between appliances. Local and remote are relative to where you are logged in. The Actifio appliance you are logged into is the local Actifio appliance and the other Actifio appliances are considered remote.

Actifio CDS and CDX Physical Appliances

Actifio CDS and CDX appliances are two-node hardware and software solutions. Actifio physical appliances are installed and configured by your Actifio representative.

Actifio Sky Virtual Appliance

An Actifio Sky appliance is a virtual machine. Actifio Sky appliances are licensed by capacity. An Actifio Sky appliance captures and manages data locally and can replicate protected data to other Actifio appliances.

Actifio Sky for AWS Cloud Appliance


An Actifio Sky for AWS appliance is a virtual machine that is licensed by capacity and resides in the AWS cloud space. An Actifio Sky appliance captures and manages application data in the AWS cloud and can replicate captured data to another location. The underlying engine for Actifio Sky for AWS is Actifio Sky.

VDisks

Actifio appliances use logical VDisks (virtual disks or volumes) to virtualize data from hosts. VDisks are taken from a pool of managed disks (MDisks) presented to an Actifio appliance from one or more internal and external arrays.

From the VDisks, the data can be deduplicated, cloned, mounted and recovered, presented for test and development work, and manipulated in other tasks. VDisks are created as needed on physical disk arrays.

There is a fixed limit of VDisks per Actifio appliance. As you create protection policies, your Actifio appliance will warn you when a configuration may exceed VDisk limits. Actifio appliances employ VDisks in slightly different ways, but the information in this section applies to all types of Actifio appliances.



VDisk Consumption

How Many VDisks Do I Have?

An Actifio appliance creates VDisks as needed from pools of MDisks on the physical disk arrays in the Actifio appliance. The applications and hosts never see the MDisks.

The VDisk limit for the Actifio CDS appliance is 5000 VDisks with the current firmware, and for the Actifio Sky appliance the VDisk limit varies with the installed capacity license (1000, 3000, or 5000 VDisks). If you have enough VDisks for your needs, but they are growing too large for your existing storage, then you must add storage. If you need more VDisks, then you need another Actifio appliance.

How Many VDisks Do I Need?

In general, each protected application or VM requires one or more VDisks for the staging disk plus the same number more VDisks per snapshot. In addition, note these rules:

- VM-level backups with a snapshot SLA consume one VDisk for each virtual disk in the VM.
- File system backups in a Windows environment consume one VDisk for each protected file system.
- File system backups in a Unix environment consume a VDisk for every 833GB protected times 1+(number of retained snapshots). You can adjust the 833GB value by changing Staging Disk Granularity in Application Advanced Settings, see ***Virtualizing and Protecting Copy Data with the Application Manager***.
- Mounts, LiveClones, and Clones of non-VM applications consume VDisks.
- On Linux and Solaris systems, filesystems and Oracle databases consume one VDisk plus another for every additional 2TB data is being protected.
- Exchange and SQL databases consume one VDisk for every volume that hosts the database.
- Each snapshot of a VDisk consumes one VDisk per snapshot per protected disk.
- Snapshots show peak usage, as new snapshots are created before old snapshots are expired.
- Each VDisk protected by Dedup-Async consumes one additional VDisk for the DAR snapshot and two VDisks on the target appliance.
- After failover and syncback, the failback operation cleans out all the syncback and failover VDisks.

VDisks are thin-provisioned, and can grow over time, as explained in [Snapshot Pool](#) on page 6.

Storage Pools

An Actifio appliance uses pools of allocated storage to store data. The amount of space to be allocated is based on how data is managed (see [Data Capture Methods](#) on page 15), how much data is involved, the type of data, its change rate, how long it will be retained, and whether or not the data is replicated to another Actifio appliance.

Snapshot Pool

The Snapshot Pool holds the most recent copies of your captured application data. Snapshot Pools retain protected data for short-term retention. Data is instantly accessible and not deduplicated. Policies determine how long data is kept in this pool and when or if data is deduplicated and moved to another pool.

Data that is replicated from a local Actifio appliance to a remote Actifio appliance via a Production to Mirror policy will land in the in the remote Actifio appliance's Snapshot Pool.

For more on the Snapshot Pool, see [Snapshot Pool](#) on page 6.

Dedup Pool

The Dedup Pool is used to retain deduplicated copies of snapshot data and to facilitate low-bandwidth replication. You can create a policy that writes VMware VM data directly to this pool, where it bypasses the Snapshot Pool and is deduplicated directly in this pool.

Data that is replicated from a local Actifio appliance to a remote Actifio appliance via a Dedup Backup DR policy will land in the in the remote Actifio appliance's Dedup Pool.

Note: *Keep space for Dedup Pools separate from space for Snapshot Pools. Dedup operations are I/O intensive. Keeping Dedup Pools separate ensures that dedup operations will not hurt the performance of the Snapshot Pools.*

For more on the Dedup Pool, see [Dedup Pool](#) on page 11.

OnVault Pool

An OnVault Pool defines the storage that can be used by an Actifio Snapshot to OnVault Policy. Data is not deduplicated when it is sent to an Actifio OnVault Pool. Actifio OnVault Pools are used for long-term storage, not for primary data storage. See also: [Sending Snapshots to an OnVault Pool's Defined Storage](#) on page 21.

Note: *The Primary Pool, act_pri_pool000, is for the Actifio appliance use. It is not a storage pool. Do not change the Primary pool or add a second pool unless instructed by Actifio Support.*

Staging Disks for VMs and Out-of-Band Applications

When you protect a VM or an out-of-band application, copies of the selected image are put into a dedicated virtual staging disk in the Snapshot pool. The Actifio appliance creates a snapshot from the image on the staging disk, and stores the snapshot in the snapshot pool for the time specified in the SLA.

Staging disks for out-of-band backups are allocated from the snapshot pool. The VDisk is thin-provisioned so it does not consume all of its space right away. Each snapshot that is created of that staging disk also consumes some amount of snapshot pool space, depending on how many changes there are.

Staging disks are not created or used for In-Band VMware VMs and applications.

Growth of Applications

If an application grows from 1TB to 2TB, a new 2TB staging disk is created. The original 1TB disk is preserved until all snapshots that depend on it are expired.

Staging Disk Formatting

Windows staging disks up to 2TB in size are MBR formatted. Those over 2TB are GPT formatted.

An Exception for Direct-to-Dedup Protection for VMware VMs

VMware VMs that are protected direct-to-dedup do not go through a staging disk because the Actifio appliance can get changed-block information directly from the VMware layer. Hyper-V VMs and all other applications get changed-block information either via Oracle RMAN or the Actifio Connector (using an Actifio staging disk), or via an in-band snapshot.

Snapshot Pool

The Snapshot pool (sometimes referred to as the Performance Pool) holds “golden copies” of application data at the points in time specified by Service Level Agreement (SLA). Snapshot space is also used to rehydrate images from the deduplication pool. The amount of data consumed is determined by whether an existing snapshot can be used. VDIs are always thin-provisioned.

Snapshot Pool space is consumed by four different kinds of VDisk:

Staging disks: Staging disks are the Actifio golden copy of the application. Staging disks are retained for as long as an application is protected and at least one snapshot exists. See [Staging Disks](#) on page 6.

Snapshot disks: These are used to preserve the state of staging disks at specific points in time. Snapshots are retained until their expiry time, but the last snapshot will never expire unless the application is unprotected or it is explicitly expired.

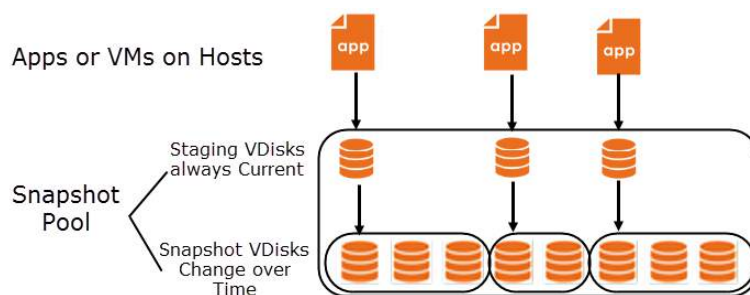
Mountable VDisk: Mountable VDIs are mountable images created at restore time from either a snapshot on a snapshot disk or a deduped image that has been rehydrated to a staging disk.

Clone disks: Clone disks are full copies of an application’s production data. Clone disks are not automatically expired.

Staging Disks

A staging disk is a VDisk created when an application is first protected. It is a copy of the production data as of the last backup invoked by the application’s SLA. Each staging disk is associated with a number of snapshots on their own snapshot VDIs. The number of snapshots for each application or VM is determined by the SLA frequency of snapshot and retention period.

Because a staging disk is a complete copy of the production application or VM, each staging disk requires as much storage space in the Snapshot Pool as the protected application or VM requires in its production storage. Snapshots made from the staging disk reference the data in the staging disk, so they are much smaller. As subsequent backups change blocks in the staging disk, the original blocks are “pushed” into the snapshot VDIs, so the snapshot appears to have constant content but contains more and more blocks over time.



Virtualized Copy Data on Staging VDIs and Snapshot VDIs in the Snapshot Pool

Understanding Snapshot Pool Consumption

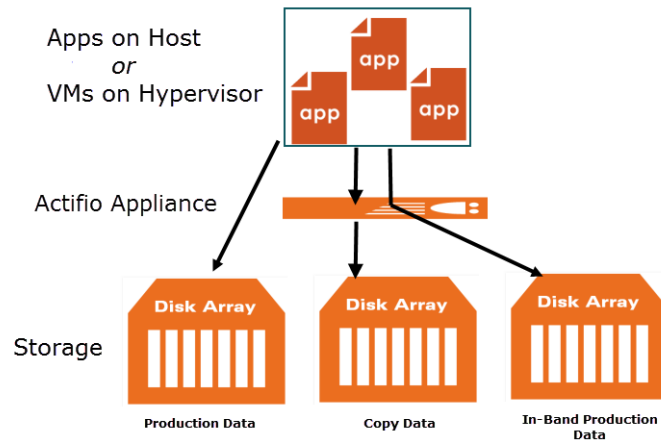
The Snapshot Pool contains both the staging disks and the snapshot disks for every protected application or VM, plus any clones and mount images that you make.

The storage space consumed in the Snapshot Pool depends upon whether the application and VM production data is in-band or out-of-band:

Out-of-Band: Production data storage is handled by a non-Actifio volume controller. Snapshots and changed-block tracking are handled by the Actifio Connector on physical hosts and Hyper-V VMs, and through VMware APIs for VMware VMs. Out-of-Band storage uses your existing storage arrays. Out-of-Band storage is described in [Out-of-Band Data and Staging Disks](#) on page 7.

In-Band: Production data storage is controlled by an Actifio CDS appliance. Snapshots and changed-block tracking are native to Actifio CDS appliance. In-Band storage is described in [In-Band Data](#) on page 9.

Note: In-Band is provided only by Actifio CDS appliances.



A Single Actifio CDS Appliance Can Protect Data In-Band and Out-of-Band

Out-of-Band Data and Staging Disks

When the Actifio appliance is not in the data path, it protects copy data by presenting a staging disk to the host. This staging disk maintains a golden copy of the application data that is protected using Actifio appliance snapshots. On Windows, application-consistent backups are made via VSS. Oracle backups on all platforms are application consistent via RMAN interfaces. Whole VM backups are application-consistent if they are configured with vmtools.

Whenever possible, change block tracking is used to minimize backup data movement. Tracking is accomplished with VSS snapshots, Oracle RMAN, and the Actifio Connector.

This section describes:

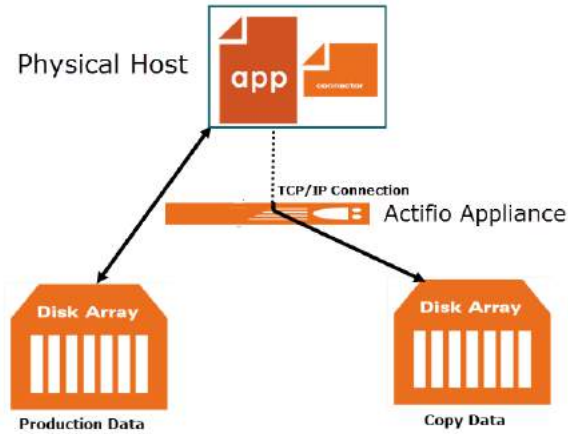
[Protecting Physical Hosts Out-of-Band](#) on page 7

[Capture Entire VMware and Hyper-V VMs](#) on page 19

[Protecting Individual Applications on a VMware or Hyper-V VM Out-of-Band](#) on page 8

Protecting Physical Hosts Out-of-Band

Applications on physical hosts are backed up with the Actifio Connector by reading data from the production data and sending this data to the snapshot pool or directly to the dedup pool. Most Windows backups are made in an application-consistent way via VSS, and database backups are application-consistent through database-specific APIs. Change block tracking is typically used to accelerate the process. Applications on non-Windows hosts are low-splash.




Protecting Applications on Physical Hosts Out-of-Band

Protecting Entire VMware and Hyper-V VMs Out-of-Band

An Actifio appliance can protect entire VMware and Hyper-V VMs Out-of-Band.

To protect entire VMware VMs, the Actifio appliance takes advantage of VMware APIs. To protect entire Hyper-V VMs, the Actifio Connector is installed on the Hyper-V server.




Protecting Entire VMs Out-of-Band

Note: An Actifio Sky appliance is a VMware VM and can be on the same ESX server as VMs it protects.

Protecting Individual Applications on a VMware or Hyper-V VM Out-of-Band

The Actifio Connector is used to protect individual applications on a VM. Once the Actifio Connector is installed on a VM, you can create policies to protect sets of applications on the VM.



Protecting Applications on a VM Out-of-Band

Note: An Actifio Sky appliance is a VMware VM and can be on the same ESX server as VMs it protects.

In-Band Data

In addition to the storage pools for copy data, the Actifio CDS appliance can also have additional storage for your applications. Actifio Optimized Storage (AOS) is controlled by the Actifio CDS appliance volume controller.

Note: In-Band is not applicable to Actifio Sky appliances.

The Actifio CDS appliance can protect a Fibre Channel presented volume using storage virtualization. In-Band means that the Fibre Channel attached back-end storage device presents volumes to the Actifio CDS appliance (using SAN LUN masking) which the Actifio CDS appliance then presents to the host over Fibre Channel or iSCSI. This storage method allows the Actifio CDS appliance to track every block of the volume and determine which blocks change over time.

When the SLA driven scheduler determines a production snapshot is required, the Actifio CDS appliance makes a copy-on-write snapshot. This action allows the Actifio CDS appliance to save space in the snapshot pool by storing only the changed blocks. These snapshots can then be used for:

- **“Instant-Backup”** - Because each snapshot does not in itself move data, copy on first write only occurs when blocks change.
- **“Instant-Restore”** - Because the Actifio CDS appliance is presenting the storage, it can simply present a previous point-in-time to the host. Any data movement needed to complete the restore is done in the background without any further host downtime.

After each snapshot, the Actifio CDS appliance can create a deduplicated and compressed image of the snapshot that is stored in the Actifio dedup pool.


Both physical hosts and VMs can be protected in-band:

[Protecting Physical Hosts In-Band](#) on page 9

[Protecting Entire VMs In-Band](#) on page 10

Protecting Physical Hosts In-Band

The Actifio Connector is used to capture applications on physical servers. The Actifio Connector achieves application-consistent backups by enabling Microsoft VSS (Shadow Copy) snapshots, and database quiescing.




Protecting Physical Server In-Band

Protecting Entire VMs In-Band

An Actifio appliance can protect entire VMware and Hyper-V VMs.

- To protect entire VMware VMs, the Actifio appliance takes advantage of VMware APIs to quiesce applications and achieve application-consistent backups.
- To protect entire Hyper-V VMs, the Actifio Connector is installed on the Hyper-V server. The Actifio Connector achieves application-consistent backups by enabling Microsoft VSS (Shadow Copy) snapshots, and database quiescing.



Protecting Entire VMware and Hyper-V VMs

Dedup Pool

The local Dedup Pool is used to retain deduplicated copies of the snapshot data for quick access, and to facilitate low-bandwidth replication.

The Dedup Pool has a warning threshold at 80% of capacity. **When this level has been exceeded, no new dedup related jobs can be run.** The Actifio appliance sends a warning when dedup pool consumption exceeds the warning level. A remote appliance will reject inbound replication jobs if its own Dedup Pool is above this threshold.

Under normal usage, deduplicated images accumulate over time, and are expired after a set retention period, having been replaced by newer images. The Garbage Collection (GC) feature identifies and reclaims the space used by the expired images. But over time, you may be capturing more or larger images, and the pool can fill up. Ultimately you may need to add more capacity, but before that becomes necessary, review this section carefully to be sure you are following best practices.

Note: Keep space for Dedup Pools on storage separate from space for Snapshot Pools. Dedup operations are I/O intensive. Keeping Dedup Pools on separate storage ensures that dedup operations will not impede the performance of the Snapshot Pools.

Understanding Dedup Pool Consumption

The Dedup Pool becomes filled with both current and expired images. GC reclaims space by sweeping away the expired images. An expired snapshot image requires as much space as it did when it was current. In general terms, expired images accumulate in the Dedup Pool at a rate determined by:

Dedup Pool Consumption = Frequency of Snapshots x Snapshot Size

Snapshot Size = Proportion of Changed Blocks (change rate) x Application Size

A 10TB application with a 50% change rate actually generates snapshot images that are 5x bigger than a 100TB application with just a 1% change rate. But if it is protected 5x as often, then they consume the same amount of the Dedup Pool.

Example of Dedup Pool Consumption

Application RPO	Snapshot Size	SLA	Dedup Pool Consumption
Small, Dynamic, Non-Critical Application	50% x 10TB = 5TB	Weekly, Friday, 6pm	5TB/week
Large, Slow-Changing, Critical Application	1% x 100TB = 1TB	Monday, Tuesday, Wednesday, Thursday, and Friday, 6pm	5TB/week

Note: This is an oversimplification of a complex subject, but the principles illustrated here can help you to understand how to keep your Dedup Pool at top performance.

Image Retention and Garbage Collection

Those images add up. An application that contributes 30GB per day to the Dedup Pool produces after 30 days close to 1TB of net new data. If these images are retained for 12 months, then after that time it has consumed 12TB of Dedup Pool space. Long retention times can quickly fill the Dedup Pool.

Image retention is how long the image is retained before it is expired. Expired images still consume Dedup Pool, so if those images are retained for only one month, but GC is never run, then the expired images still accumulate.

To reclaim that space, the unneeded images must be expired and then Garbage Collection (GC) must be run. Dedup Pool Consumption grows when the rate of incoming images exceeds the rate at which they are swept out by GC.

To reduce the amount of dedup pool space consumed:


- Retain images no longer than needed, and set GC accordingly.
- Use a mixed retention policy to decrease the number of unnecessary images. For example, retain daily images for one month, weekly images for three months, and monthly images for one year.
- Expiring images or reducing retention rates will not have an immediate affect on the dedup pool. You must run garbage collection and sweep to reclaim the space.
- Run Garbage Collection (GC) and sweep on a regular basis. Once per month is common.

How Dedup Pool Space is Consumed and Reclaimed: Garbage Collection

Space in the Dedup Pool is consumed and reclaimed over time. If the Dedup Pool becomes too full, then there is insufficient space to process new snapshots, and jobs start to fail. Knowing how and when to reclaim space is important for best operation.


The Simple View

If you had an ideal system with snapshots all the same size being deduplicated and expired at even intervals, you might expect the Dedup Pool consumption over time to look like this:



An Oversimplified and Incorrect View of Dedup Pool Consumption: Snapshots and Expiry


But expiring an image is not enough to reclaim the space. The space is only reclaimed after a Garbage Collection run has enumerated the blocks to keep and then reclaimed the remaining space. So the expired images accumulate until the GC finishes:



A Simplified, More Correct View of Dedup Pool Consumption: GC Reclaims Space

The Detailed View, for When Your Dedup Pool is Filling Up

When GC runs regularly and normally, the image above is enough to understand operations. But the process is a little more complex, and the difference is important when your Dedup Pool is approaching capacity. This is because GC has two phases, the GC Enumeration Phase and the GC Sweep Phase.



GC Increases Dedup Pool Consumption Before Reclaiming Space

If your Dedup Pool is approaching capacity, a GC run is very important, but the GC Enumeration Phase may temporarily increase your pool consumption to the point that jobs start to fail. It is important to monitor Dedup Pool consumption and to run GC monthly during times of low activity. With regard to keeping your Dedup Pool from filling up, you need to know these three things:

1. During the GC Enumeration Phase, Dedup Pool consumption continues to increase.
2. The GC Enumeration Phase can take a long time to run, even days, and it must not be interrupted. The Sweep Phase is usually faster, and you can set it to run at intervals such as overnight.
3. If necessary, the GC Enumeration Phase and the GC Sweep Phase can be run at different times.

The Impact of Garbage Collection on Dedup Performance

GC is resource-intensive and time-consuming, sometimes taking days to complete. Initial ingests of new applications cannot run while GC is running.

GC consumes system resources. All jobs except initial ingests continue to run as scheduled, but some of them take longer to complete due to resource competition from GC. A long-running GC can result in a backlog of jobs and even SLA violations from jobs that do not complete within their scheduled time.

GC runs in the Dedup Pool, so Dedup, Direct-to-Dedup, and Remote Dedup jobs are heavily impacted. It is a good idea to reschedule these if a GC is planned, or is running overlong. Snapshot (protection) jobs are unaffected.

A good rule of thumb is to identify the applications that require the most Dedup Pool consumption (see [Understanding Dedup Pool Consumption](#) on page 11), and then set GC to run at four times the retention of those images; for example if those applications are typically retained for one week, then run GC monthly.

Policy Templates and Policies

A Policy Template is a collection of policies. A policy defines:

- The source of the data managed by the policy
- Type of the protection operation
- Frequency of the protection operation
- How long to retain the data
- Whether data is replicated

Multiple policies within a template allow you to create a single template that defines the short term and long term retention of data as well as whether data will be replicated and how long replicated data will be retained.

Depending on the Actifio appliance type, a Policy Template can be made up of one or more of the following policies:

Production to Snapshot defines when and how often production data will be captured and how many snapshots are retained. Data recovery from the Snapshot Pool is fast because images are not deduplicated and are stored in the local Actifio appliance's Snapshot Pool. Snapshots are meant for short term retention. See [Production to Snapshot Policies](#) on page 31 for details.

Snapshot to Dedup Backup defines when to deduplicate snapshot data and how long to retain the deduplicated data. Data in the Dedup Backup Pool is meant for longer term retention. See [Snapshot to Dedup Backup Policies](#) on page 31 for details.

Production to Mirror defines how data will be replicated to a Mirror Pool (a Snapshot Pool on a remote Actifio appliance). Data in the Mirror Pool is meant for instant recovery in a disaster recovery scenario. See [Snapshot to Dedup Backup Policies](#) on page 31 for details.

When creating a Production to Mirror Policy you have the choice of replicating data via:

- o StreamSnap - Used in environments that have high bandwidth networks
- o Dedup-Async Replication (DAR) - Used in environments where network bandwidth is constrained

For details on replicating data, see [Replicating Data](#) on page 20.

Dedup Backup to Dedup DR defines when to replicate deduplicated data to a remote Actifio appliance's Local Dedup Pool and how long to retain the data in that pool. Data in the Dedup DR Pool is meant for retention of data in case of a disaster at the local appliance's site. See [Dedup Backup to Dedup DR Policies](#) on page 32 for details.

Production Direct-to-Dedup defines when to deduplicate VMware VMs directly from production data and how long to retain the deduplicated data. Capturing VMware VMs directly to a Dedup Backup Pool is meant for long term retention when instant access from a Snapshot Pool is not required. See [Production Direct to Dedup Policies](#) on page 34 for details.

Snapshot to OnVault defines when to send Production to Snapshot data to the storage defined by an Actifio OnVault Pool and how long to retain the data. Snapshot to OnVault Policies are meant for long-term retention of data. See [Snapshot to OnVault Policies](#) on page 32 for details.

Policy Templates are:

- Created in the Actifio Service Level Architect (SLA) service.
- Applied to applications in the Application Manager service.

Resource Profiles

Resource profiles define where application data is retained. They define which pool to use; Snapshot, Dedup, or OnVault. Pools specified in Resource profiles are used along with policy templates to form an SLA for an application. Resource Profiles are:

- Created in the Actifio Service Level Architect (SLA) service.
- Applied to applications in the Application Manager service.

Data Capture Methods

An Actifio appliance provides three methods for managing production data: Out-of-Band, In-Band, and LAN-Free. A single Actifio appliance can capture some applications Out-of-Band, some In-Band, and other applications LAN-Free.

Out-of-Band

This is the most common method used when managing data. Production data is controlled by a non-Actifio storage controller on your existing storage arrays. The Actifio appliance operates outside of the application's data path and leverages the IP network. The Actifio appliance moves and captures the application data separately from where the application writes its primary storage.

Snapshots of application data are captured and stored on a staging disk presented to the application host via Fibre Channel or iSCSI.

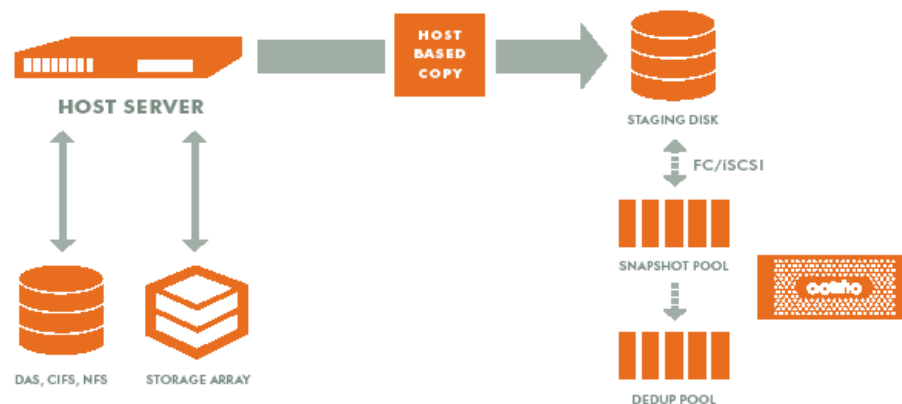
The Out-of-Band method will meet the needs of most users who want to capture and access:

Hypervisors: VMware and Hyper-V

Applications: Oracle, SQL, Exchange, SharePoint, SAP on Oracle

File Systems: Windows, Unix, Linux file systems

An Actifio appliance captures copy data by presenting a staging disk that maintains a golden copy of the application's historical data. Data is captured using Actifio snapshots



Typical Out-of-Band Capture Method

When capturing data Out-of-Band:

- A staging disk is mounted on a server via Fibre Channel or iSCSI.
- An initial capture of the entire image is made to the staging disk.
- Subsequent captures consist only of incremental changes. This is made possible by taking advantage of change block tracking via the Actifio Connector or by VMware APIs. See [Capture Mechanisms](#) on page 17 for details.
- The staging disk is unmounted on the server.
- A snapshot of the staging disk is made on the Actifio appliance.

In-Band

Note: In-Band is not applicable to Actifio Sky or CDX appliances.

With the In-Band method, production data storage is controlled by the Actifio appliance. Snapshots and changed-block tracking are native to the Actifio appliance. The appliance resides in the data path between the SAN and the application host.

You can use the In-Band method on production data that is on a Fibre Channel LUN(s) on an Actifio appliance-supported storage array AND if one or more of the following conditions apply:

- The application is not a supported Out-of-Band application; for example, DB2 or a custom application.
- The local RPO requirements are shorter than what is practical for Out-of-Band. For example, when snapshots are required every 15 minutes.
- The remote RPO requirements are shorter than what Actifio Dedup-Async Replication (DAR) allows. For example, if a requirement states instant sync/async.
- There is a large amount of data; for example, a 10TB database
- The applications and the files on them need to be captured. It is more efficient to capture blocks of data than the applications and their files. For example, a Linux file system with 21 million files.

LAN-Free

To capture VMware VMs, an Actifio appliance can employ LAN-Free data capture.

Note: Capturing VMs does not apply to Actifio Sky for AWS.

To use the LAN-Free data capture method the SAN administrator simply has to:

1. Use Fibre Channel SAN zoning to provide the Actifio appliance access to the storage controller that manages the ESXi datastores.
2. Ensure that the storage controller is supported by the Actifio appliance.
3. Define the Actifio appliance as a host
4. LUN mask all datastore LUNs to the Actifio appliance host.

The Actifio appliance will detect whether the ESXi datastore LUNs are accessible via Fibre Channel. If they are available, data will be moved across the Fibre Channel SAN.

In all other aspects LAN-Free is the same as Out-of-Band.

Note: If the SAN administrator fails to map the required datastores, or maps a required datastore away from the Actifio appliance, then the appliance will switch to LAN-based data capture.

LAN-Free vs Out-of-Band

LAN-Free and Out-of-Band Fibre Channel are both LAN-Free, although Command and Control is still done over the LAN. The difference is the source HBA.

- In LAN-Free, Actifio reads the data direct from the storage HBA (Actifio is initiator)
- In Out-of-Band over Fibre Channel, Actifio is sent the data by the host HBA (Actifio is target).

Capture Mechanisms

An Actifio 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 and capture the changes that occur between capture operations. To track those changes the Actifio appliance uses either the Actifio Connector or VMware API calls.

Note: Capturing VMs does not apply to Actifio Sky for AWS.

The Actifio Connector

The Actifio Connector is used to capture selected applications and for capturing entire Hyper-V VMs. The Actifio Connector is a small-footprint, operating system specific, lightweight service that can be installed on either virtual or physical servers.

The Actifio Connector provides a more granular data capture capability than what is provided by VMware API calls. It allows you to: Capture selected applications, capture applications that cannot be snapped by VMware, capture Microsoft® SQL Server clusters, and Microsoft Exchange Database Availability Groups (DAGs).

Specifically, Actifio Connectors:

- Discover applications
- Quiesce applications, for application consistency during capture
- Enables change block tracking for Actifio's incremental forever capture strategy
- Capture and manage transaction logs:
 - o Capture database(s) and logs with one policy template
 - o Truncate database transaction logs
 - o Roll database transaction logs forward for point-in-time recovery when accessing virtual copies
- If multiple applications are resident on a server, a single policy template can be applied to multiple applications.
- Avoids VMware VMs "stun" issues

For Hyper-V Servers, the Actifio Connector also enables the capture of entire Hyper-V VMs.

The Actifio Connector also enables scripting on the hosts on which it is installed. The Actifio Connector (host side) scripts can be invoked for:

- On-demand jobs triggered from the Actifio CLI with the **-scripts** argument.
- Pre- and post-phases of an Actifio appliance Workflow job.

For detailed step-by-step instructions on how use Actifio appliance scripting, see **Connecting Hosts to Actifio Appliances** in the Actifio Documentation Library or the ActifioNOW customer portal.

VMware API Calls

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

- Enable change block tracking -for Actifio's incremental forever capture strategy.
- Quiesce applications for application consistency during capture.

When an entire virtual server is captured, 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 Actifio appliance directly and then optionally migrated to a new, permanent location.

Virtual machines and their applications can be grouped and captured with a single Policy Template.

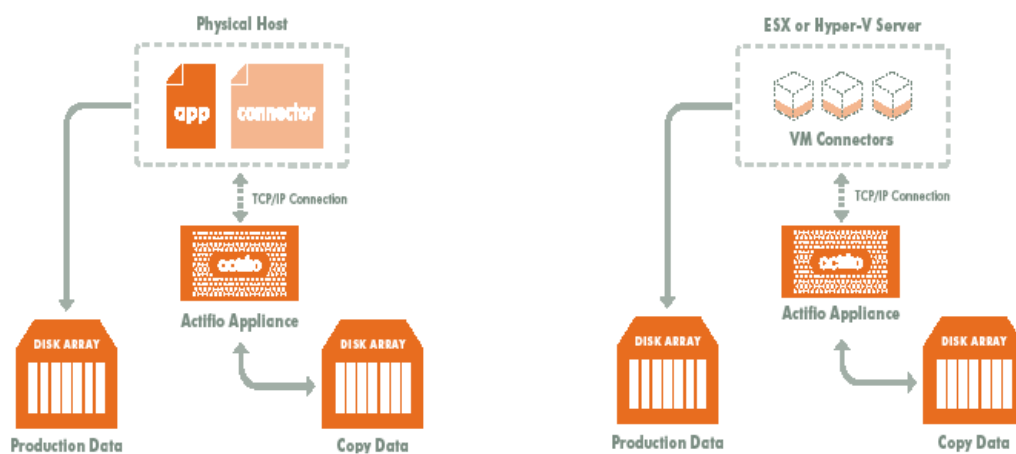
Capture Options

An Actifio appliance allows you to:

- [Capture Individual and Groups of Applications](#) on page 18
- [Capture Application Data in Actifio Consistency Groups](#) on page 18
- [Capture a VM's Applications and Boot Volume](#) on page 19
- [Capture Entire VMware and Hyper-V VMs](#) on page 19
- [Capture Generic Applications](#) on page 19
- [Capture Database Logs](#) on page 20

Capture Individual and Groups of Applications

The Actifio Connector is used to capture individual and groups of applications on physical and virtual servers.



Managing Individual or Groups of Applications

Installing the Actifio Connector on a physical server or VM allows you to create a single Policy Template to capture all applications on the server or several Policy Templates to capture groups of applications.

Capture Application Data in Actifio Consistency Groups

A consistency group is enabled by the Actifio Connector. As the name implies, consistency groups ensure consistent point-in-time capture and recovery across multiple applications on the same host.

To achieve application consistency, members of a consistency group are quiesced and captured together via a single Policy Template.

If Actifio's Database Log Backup option (Microsoft SQL Server and Oracle only) is enabled on a Snapshot policy, then all databases captured by the Policy Template in which the Snapshot policy resides can be recovered to the same point-in-time. Recovery and rolling forward of the logs (for databases) in a group is performed via the Actifio user interface with a single action.

In addition to making capture and recovery operations easy and fast, consistency groups consume fewer system resources (VDisks).

Capture a VM's Applications and Boot Volume

Note: Capturing VM Boot Volumes does not apply to Actifio Sky for AWS.

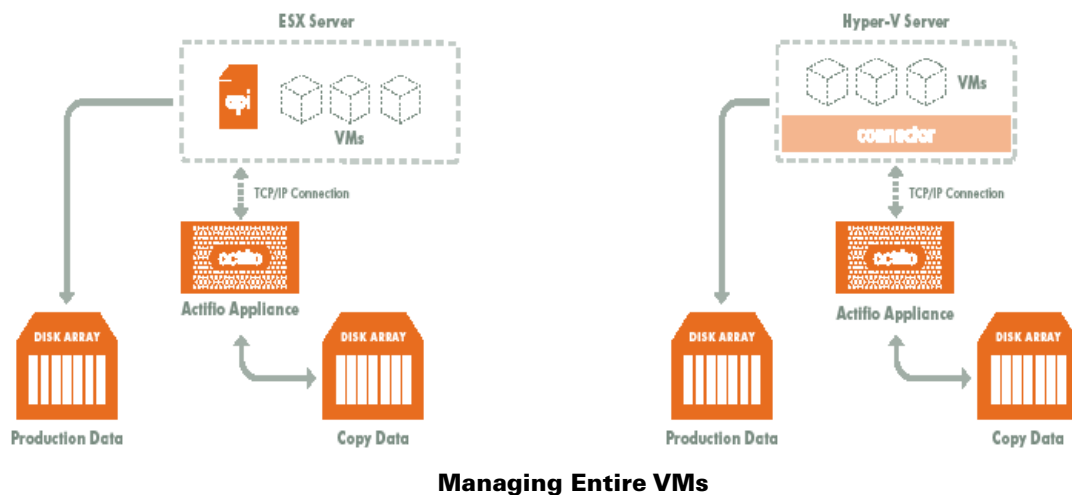
When managing applications on VMs you have the option of also capturing the VM's boot volume.

When a VM's boot volume is captured along with its applications, an image can be presented that is a fully functional VM. The image can then be migrated to a new, permanent location if needed.

Capture Entire VMware and Hyper-V VMs

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

Note: Capturing VMs does not apply to Actifio Sky for AWS.



Note: An Actifio Sky appliance is a VMware VM that can be on the same ESX server as the VMs it captures.

When an entire virtual server is captured, 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, it can be migrated to a new, permanent location if needed.

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

Capture Generic Applications

The Generic Application type is used to capture data for applications that do not have dedicated, application specific, Actifio data capture processes. In addition, there are some instances where applications that do have dedicated processes will use the generic application type for data capture.

Capture Database Logs

Database log capture is enabled in a Snapshot policy's Advanced Options. It enables a single Snapshot policy to capture logs for Microsoft SQL Server databases, Oracle databases, and consistency groups that contain Microsoft SQL Server databases or Oracle databases.

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 capture frequency is every 24 hours, the log file capture frequency must be less than every 24 hours.

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. For example, if a database's Snapshot data is kept for three days and its Dedup data kept for seven days, you can define log retention to span all seven days.

In this example, a single captured database image can be selected and its logs can be rolled forwards over the seven day period.

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

If required, you can specify to replicate Oracle or Microsoft® SQL Server database transaction logs to a remote Actifio appliance. You can use the logs at the remote site for any database image within the retention range of the replicated logs. Log replication uses StreamSnap technology to perform the replication between the local and remote Actifio appliances; the replication goes directly from the snapshot pool of the local Actifio appliance to the remote snapshot pool. For log replication jobs to run, there must be a replication policy (StreamSnap, Dedup-Async, or Remote Dedup) included in the template and a resource profile that includes a remote Actifio appliance, and at least one successful replication of the database must first be completed.

Replicating Data

Replication of copy data to remote storage protects the data in the event of disaster at the primary site and reduces the amount of storage required at the primary site. The goal of replication is to get your data back in situations of data loss and impact to your production systems due to issues such as a hardware failure, software issues, or a site event. Data replication also supports the creation of remote copies of Test/Dev, QA, and Analytics data. Data can be replicated from one Actifio appliance to a second (remote) appliance or to the cloud for recovery, disaster recovery, or test/development purposes.

Note: Details about the different types of replication methods supported by the Actifio appliances along with step-by-step instructions on their use can be found in **Replicating Data Using Actifio Appliances** in your Actifio Documentation Library and on the ActifioNOW customer portal.

Your SLA templates determine the method, schedule, and frequency of how data replication to a remote site is to be performed. The SLA template defines how to move and store data efficiently to the remote Actifio appliance. Data replication is controlled by the individual template policies:

- Dedup Backup policies use an Actifio proprietary replication engine to replicate data to a second Actifio appliance. Dedup Backup is efficient for the long-term storage (weeks to months) of captured and deduplicated data to a remote Actifio appliance. Dedup Backup is intended to retain data for a medium to long-term retention period (for example, 3 months to a year). In addition, Second Hop Replication policies allow you to replicate dedup backup data to a second location operating as the second-leg in a multi-hop configuration of joined Actifio appliances. For details see [Dedup Backup to Dedup DR Replication](#) on page 21.

- Production to Mirror policies protect your application or VM data against a site failure by having a full copy of that data mirrored to a remote production site. Applications are kept up-to-date and can be re-started at a moment's notice at the remote site by accessing data from the remote DR copy. Data mirroring can be considered as access optimized replication to a remote site. For details see [Production to Mirror Policy Replication](#) on page 21.
- Snapshot to OnVault policies use an HTTPS connection to send data to storage defined by an Actifio OnVault Pool. Data sent to storage defined by an Actifio OnVault Pool is not deduplicated. The compress option is on by default in Actifio OnVault Pools. For details see [Sending Snapshots to an OnVault Pool's Defined Storage](#) on page 21

Dedup Backup to Dedup DR Replication

Dedup Backup replication is efficient for long-term storage of captured and deduplicated data to a remote Actifio appliance. Data replicated using a Dedup Backup policy is transmitted from the local Actifio appliance dedup pool to the dedup pool managed by another Actifio appliance. The need for and number of Dedup DR copies to retain on a second Actifio appliance for long-term data recovery (LTDR) is driven by offsite retention requirements for the data.

Dedup Backup uses a Dedup Backup to Dedup DR policy. Dedup Backup replication is incremental, globally deduplicated, and compressed and encrypted in flight. The Dedup Backup replication process begins after the deduplication process completes. A proprietary deduplication-aware replication protocol enables the transmission of only the globally unique blocks, which minimizes the bandwidth required to move data between Actifio appliances.

Blocks are compressed and encrypted in flight for the most efficient bandwidth utilization. Deduplication and compression optimize the data set for transport between sites, eliminating the cost of WAN optimization.

Dedup Backup replication also provides an added benefit of allowing data to be replicated to a remote site, and then from that remote site to a second remote site. This feature is referred to as multi-hop. You can use this process if need to perform on-demand replication of a backup dedup image to a remote Actifio appliance that is part of a multi-hop configuration of joined Actifio appliances.

Sending Snapshots to an OnVault Pool's Defined Storage

The Snapshot to OnVault policy allows you to send snapshot data to a location defined by an Actifio OnVault Pool. A schedule within the policy is used to send the most recent snapshot taken by the Policy Template's Production to Snapshot policy to the location defined by the Actifio OnVault Pool. Actifio OnVault Pool storage is typically used for long-term retention. For details on the OnVault Pool, see [OnVault Pool](#) on page 5.

When sending data to a storage defined by an Actifio OnVault Pool, an HTTPS connection is used to ensure data security over the network. The OnVault Pool's compression option is on by default to minimize network traffic.

Data sent to the Actifio OnVault Pool storage is not deduplicated. However, after the initial ingest of the full snapshot, only the changes to data are sent to the OnVault Pool. This is the same incremental forever model used by other Actifio policies.

When accessing data in an Actifio OnVault Pool's storage:

- Actifio CDS and Sky appliances can create clones.
- Actifio CDS and Sky appliances can mount data, but because data will first be copied to the snapshot pool then mounted, it is not recommended.
- LiveClones cannot be created.

Production to Mirror Policy Replication

Production to Mirror policies provide a means to replicate a copy of the application or VM data to a target Actifio appliance and to have data access without a restore window, providing for very low RTO. As needed, you have the ability to perform a failback to the production site with an identical set of data that is mirrored between the local and remote Actifio appliances.

StreamSnap

StreamSnap facilitates high-availability by allowing you to keep a remote copy of an application's storage and configuration 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 data snapshots that are not deduplicated to a remote Actifio 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 Actifio 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 Actifio 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.

You can retain snapshot images from multiple available points in time at the remote site by applying retention in a StreamSnap policy. When retaining snapshot images at the remote Actifio appliance, a new snapshot image will be created at the remote appliance with an expiration date determined by the policy settings. Each remote snapshot image supports all operations available with a local snapshot image when accessed from the Application Manager.

StreamSnap replication requires a reliable network connection to replicate data snapshots to the remote Actifio appliance. The bandwidth required on the network connection is directly related to the application size (initial copy) and amount of change (for incremental updates).

Dedup-Async (DAR)

Dedup-Async Replication (DAR) allows you to keep a remote copy of an application's data up-to-date and ready to be used in a failover scenario, facilitating high-availability and redundancy. When a DAR-managed application fails for any reason, you can mount a failover image of the application at the remote site. When the problem has been resolved, then you can incrementally reverse replicate the changes made at the DR site to the primary site and then failback the application to the production site.

DAR is an Actifio-proprietary form of replication where initially a full copy of the application data is replicated to a target Snapshot pool on the second Actifio appliance or when a VM is replicated to the datastores of an ESX server. Dedup-Async replication sends deduplicated and compressed data over the network at a fraction of the bandwidth required for traditional replication technologies.

Once the Dedup-Async job takes a snapshot, it deduplicates the data, replicates the deduplicated data to another Actifio appliance, rehydrates that data on the second Actifio appliance, and updates the full copy of data on the second Actifio appliance to provide the flexibility of instant access at the remote location. This ensures that a full, up-to-date copy of data is ready and available on the second Actifio site.

Because the data is deduplicated before it is replicated, DAR requires less network bandwidth than StreamSnap replication but it does require additional Actifio system resources.

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.*

Actifio NAS Director

The Actifio NAS Director is an addition for your Actifio appliance that turns the appliance into an efficient platform for reliably capturing, archiving, replicating, and recovering a variety of unstructured data from large file systems. Petabyte-sized unstructured datasets have become common in the enterprise. Actifio NAS Director enables you to capture the file system production data that resides on the large-scale NAS servers in your environment. Each Actifio NAS Director node provides the necessary CPU, memory, and network bandwidth for the Actifio appliance to capture and export NAS file-level data, and to dedup this data. The Actifio appliance supports network connectivity with up to eight Actifio NAS Director nodes.

Partial or full NAS shares protected by the Actifio n appliance using a Actifio NAS Director Node are called NAS datasets. NAS datasets are the data from the Isilon and NetApp NAS servers that are virtualized by the Actifio appliance and stored on an Actifio NAS Director node. You protect a NAS dataset by binding a resource profile and a template to it to make an SLA. You can then mount a backed-up NAS dataset and export it to an NFS or CIFS share.

Note: NAS datasets, NAS servers, and the Actifio NAS Director node are not supported on Actifio Sky for AWS.

Managed Data License

Managed Data License (MDL) is a measure of Actifio-protected copy data. The measure of protected data varies by application. For example:

- File Systems: MDL equals the total size of protected files.
 - Oracle databases: the total size of protected database files.
 - SQL Server, Exchange, and SharePoint: the total size of protected database files, including any log files OR the size of the disks (if it is inband).
- Generic applications: the total size of all protected volumes.

3 Capturing Applications Overview

This chapter presents high-level descriptions of the processes used to capture an application:

[Adding Hosts](#) on page 26

[Discovering Applications and VMs](#) on page 27

[Creating Policy Templates](#) on page 28

[Creating Policies](#) on page 30

[Resource Profiles](#) on page 35


[Applying Policy Templates and Resource Profiles](#) on page 36

Detailed, application-specific instructions on how to capture an application can be found in the Actifio Documentation Library:

1. Add hosts that host applications using the Domain Manager service. For detailed, application-specific instructions see the **Connecting Hosts to Actifio Appliances** guide in the Actifio Documentation Library or the ActifioNOW customer portal.
2. Create Actifio appliance policy templates and resource profiles using the SLA Architect service. Policies and templates define when and where to capture data and how long to keep it. For detailed, application-specific instructions see the **Planning and Developing Service Level Agreements** guide in the Actifio Documentation Library or the ActifioNOW customer portal.
3. Discover applications and VMs using the Application Manager service. For detailed, application-specific instructions see the **Virtualizing and Protecting Copy Data with the Application Manager** guide in the Actifio Documentation Library or the ActifioNOW customer portal.
4. Assign an Actifio Policy Template and resource profile to an application using the Application Manager service. For detailed, application-specific instructions see the **Virtualizing and Protecting Copy Data with the Application Manager** guide in the Actifio Documentation Library or the ActifioNOW customer portal.


Adding Hosts

The first step in capturing an application is to add the host on which the application(s) reside. To add a host, open the Domain Manager service and click the service menu icon:



Domain Manager's Service Menu

From the service menu, select **New Host**. A page is displayed that allows you to define the host to be captured.



Add New Host

From this page you can add servers that host applications and hypervisors that host VMs. Individual VMs are discovered via the Application Manager service.


Note: Adding VMs does not apply to Actifio Sky for AWS.

Discovering Applications and VMs

After a host is added, use the Application Manager service to discover applications on physical servers, VMs on hypervisors, and applications on VMs.


Note: *Discovering VMs does not apply to Actifio Sky for AWS.*

From the Application Manager's Navigation Pane, select a discovered physical server, hypervisor, or VM. From the Application Manager's service menu, select either Discover APP(s) or Discover VM(s):



Application Manager's Service Menu

When you select Discover APP(S) or Discover VM(s), you are prompted to select which of the discovered hosts or hypervisors you want to discover applications or VMs.



Select Discovered Applications

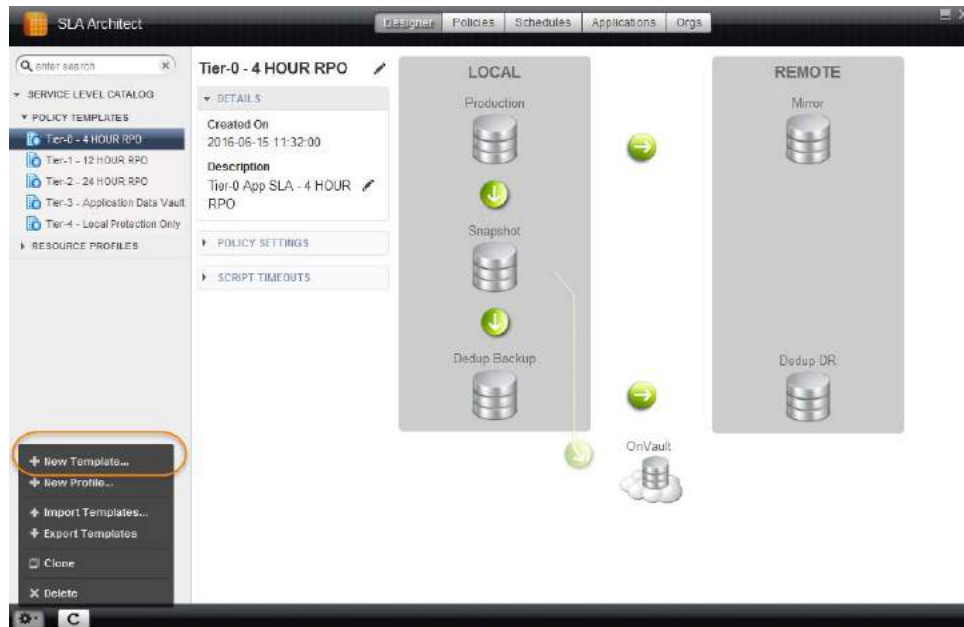
Creating Policy Templates

Policy templates are made up of one or more policies. Policy templates provide a high-level wrapper for an end-to-end definition of capturing application data. For example, if you need to capture an image as a local snapshot and replicate that image off to another Actifio appliance, the policy template will contain both the local snapshot and remote off site policies.

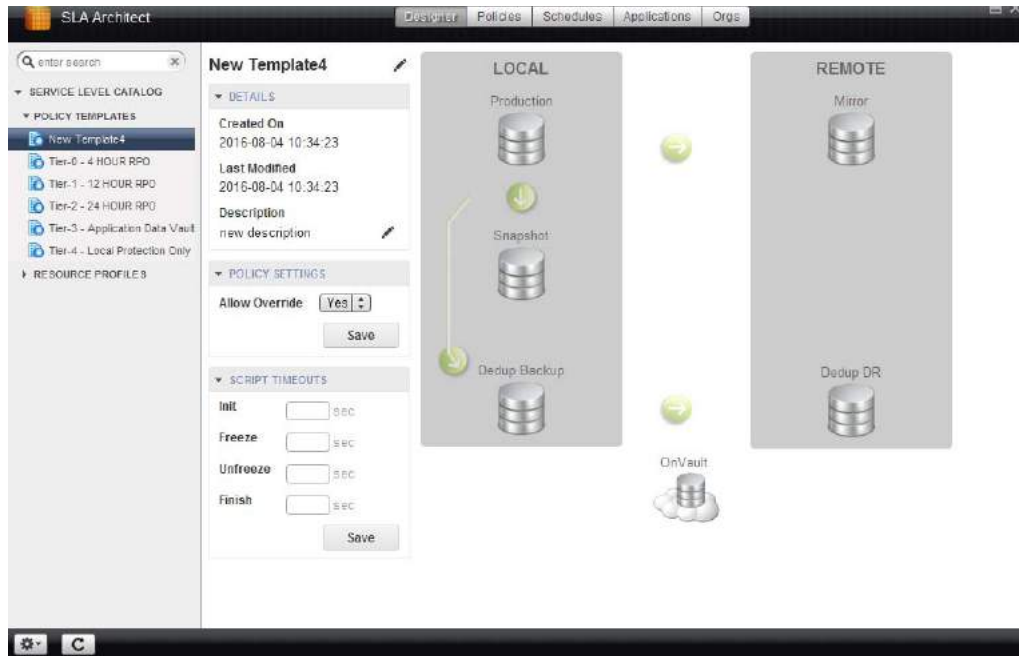
Note: See *Planning and Developing Service Level Agreements* in the Actifio Documentation Library for detailed best practices for configuring policies in Policy Templates.

To create a new template:

1. From the SLA Architect service menu, click **New Template...**



The New Template page is displayed:



New Template Page with Policy Settings Shown


2. The default name of the template is New Template. Click the New Template name above the Created on Date and rename the template as needed.
3. The policy settings allow the application-specific Advanced Settings parameters that have been defined for the individual applications in the Application Manager to override the Advanced Settings defined in the SLA template's policies.

Once the policy template is created, create the individual policies that comprise the SLA template.

Creating Policies

Policies define how often to capture an application, how long to retain the captured application and when applicable, where and how to replicate the captured application's data.

The green arrows in the SLA Architect represent the policies within a template that control data flow to the various pools.



Green Arrows Represent Policies In a Policy Template

A Policy allows you define whether its schedule will run:

- **Within a Window** - A period of time in which jobs are allowed to start.
- **Continuous** - Defines when its first job can start but as the name implies, allows subsequent jobs to run at a frequency without regard to any time boundary.

Where applicable, SLA Template Policies allow you to define the rules for determining whether or not a data protected by a policy meets your requirements. The Actifio appliance automatically calculates and sets default SLA Compliance settings. Default settings are based on whether the policy is set to windowed or continuous, the policy type, and Actifio-recommended best practices. The default settings calculated will meet the needs of most users.

If data is being protected according to your needs, then it is considered to be in compliance.

The SLA Templates are made up of the following types of policies

- [Production to Snapshot Policies](#) on page 31
- [Snapshot to Dedup Backup Policies](#) on page 31
- [Dedup Backup to Dedup DR Policies](#) on page 32
- [Snapshot to OnVault Policies](#) on page 32
- [Production to Mirror Policies](#) on page 33
- [Production Direct to Dedup Policies](#) on page 34

Note: See *Planning and Developing Service Level Agreements* in the Actifio Documentation Library for detailed best practices for configuring policies an SLA template.

Production to Snapshot Policies

Production to Snapshot policies define when and how often data will be written to the Snapshot Pool and how long that data will be retained.

In addition, application-specific advanced settings can be applied to the policy. If the template in which this policy resides allows overrides. When you apply a policy template to an application in the Application Manager service, you will be given the opportunity to override the advanced settings defined in this policy.

To create a Production to Snapshot policy:

1. Select a template from the left-hand Policy Template list.
2. Click the green arrow between **Production** and **Snapshot**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.




4. Specify the schedule type for the policy: Windowed or Continuous. The default is Windowed.
5. Enter policy schedule and retention settings that meet your RTOs and RPOs.
6. Set **Configure SLA compliance settings**. In most cases, the default settings will meet your needs.
7. Set the advanced settings that apply to the application you are capturing.
8. Save the settings.

Snapshot to Dedup Backup Policies

Snapshot to Dedup Backup policies define when data from the Snapshot Pool will be deduplicated and how long the deduplicated data will be retained in the Local Dedup pool.

To create a Snapshot to Dedup Backup policy:

1. Select a template from the left-hand Policy Template list.
2. Click the green arrow between **Snapshot** and **Dedup Backup**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.



4. Specify the schedule type for the policy: Windowed or Continuous. The default is Windowed.
5. Enter policy settings that meet your RPOs and RTOs.
6. Set **Configure SLA compliance settings**. In most cases, the default settings will meet your needs.
7. Save the settings.

Dedup Backup to Dedup DR Policies

The Dedup Backup to Dedup DR policy defines when deduplicated data from the Local Dedup pool will be replicated to a Local Dedup pool on a remote site. It also defines how long the replicated data will be kept on the remote site.

If no other policies are defined in the template, then this policy can be used as the second hop in a multi-hop replication scheme.

To create a Dedup Backup to Dedup DR policy:

1. Select a template from the left-hand Policy Template list.
2. Click the green arrow between **Dedup Backup** and **Dedup DR**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.



4. Specify the schedule type for the policy: Windowed or Continuous. The default is Windowed.
5. Enter policy settings that meet the requirements of your RPOs and RTOs.
6. Set **Configure SLA compliance settings**. In most cases, the default settings will meet your needs.
7. Save the settings.

Snapshot to OnVault Policies

The Snapshot to OnVault policy defines when Production to Snapshot data will be replicated to an Actifio OnVault Pool. It also defines how long the replicated data will be kept in an Actifio OnVault Pool.

To create a Snapshot to OnVault Policy:

1. Select a template from the left-hand Policy Template list.

Note: For new policy templates, you must first create a Production to Snapshot Policy. Once the Production to Snapshot Policy is created, the “green arrow” between the snapshot and Actifio OnVault will be activated.

2. Click the green arrow between **Snapshot** and **OnVault**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.



4. Enter policy settings that meet the requirements of your RPOs and RTOs.
5. Save the settings.

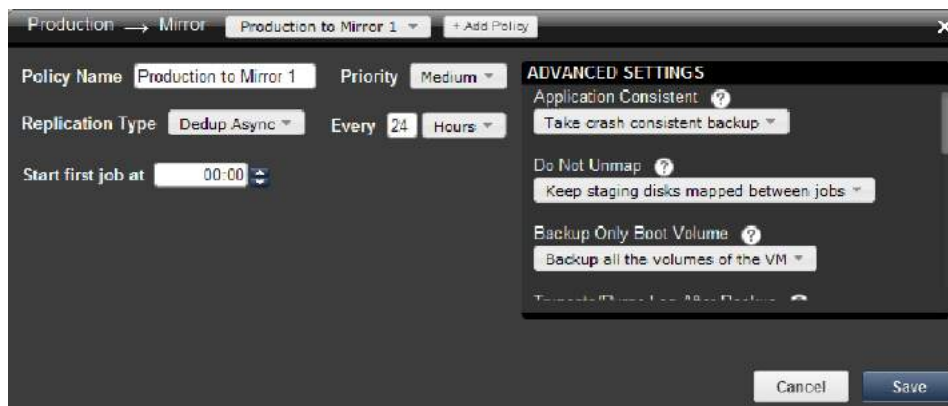
Production to Mirror Policies

The Production to Mirror policy defines how data will be replicated to a Snapshot Pool on a remote site (Dedup-Async or StreamSnap) and how often replication should occur.

In addition, application specific advanced settings can be applied to the policy. If the template in which this policy resides allows overrides: When you apply a policy template to an application in the Application Manager, you will be given the opportunity to override the advanced settings defined in this policy.

To create a Production to Mirror policy:

1. Select a template from the left-hand Policy Template list.
2. Click the green arrow between **Local Production** and **Remote Mirror**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.



4. Select the replication type (for example, **Dedup-Async** or **StreamSnap**).
5. Enter policy settings that meet your off site RPOs and RTOs.
6. Save the settings.


Production Direct to Dedup Policies

Production Direct to Dedup policies are used to capture VMware VMs directly to a dedup pool. They define when to deduplicate VMware VMs and how long to retain the deduplicated data. VMware VMs can be deduplicated directly because the Actifio appliance can get changed-block information directly from VMware APIs.

In addition, VMware specific advanced settings can be applied to the policy. If the template in which this policy resides allows overrides, then when you apply a policy template to an application in the Application Manager, you will be given the opportunity to override the advanced settings defined in this policy.

To create a Production Direct to Dedup policy:

1. Create a new policy template.
2. Click the green arrow that points from **Production to Dedup Backup**. The policy window appears.
3. Click the **+(Plus)** icon and the policy settings are displayed.




4. Specify the schedule type for the policy: Windowed or Continuous. The default is Windowed.
5. Enter policy settings that meet the requirements of your RPOs and RTOs.
6. Set **Configure SLA compliance settings**. In most cases, the default settings will meet your needs.
7. Save the settings.

Resource Profiles

A resource profile specifies where data will be backed up and the Snapshot Pool to be used. If you will replicate data in the Snapshot Pool to a Mirror Pool (Snapshot Pool) on a remote site, you will need to select the remote site from the Remote drop down list. Only those remote sites that have exchanged certificates with and joined the local site will be listed in the Remote drop down list.

To create a resource profile:

1. From the SLA Architect user interface's service menu, click **New Profile...**



The New Profile page is displayed:



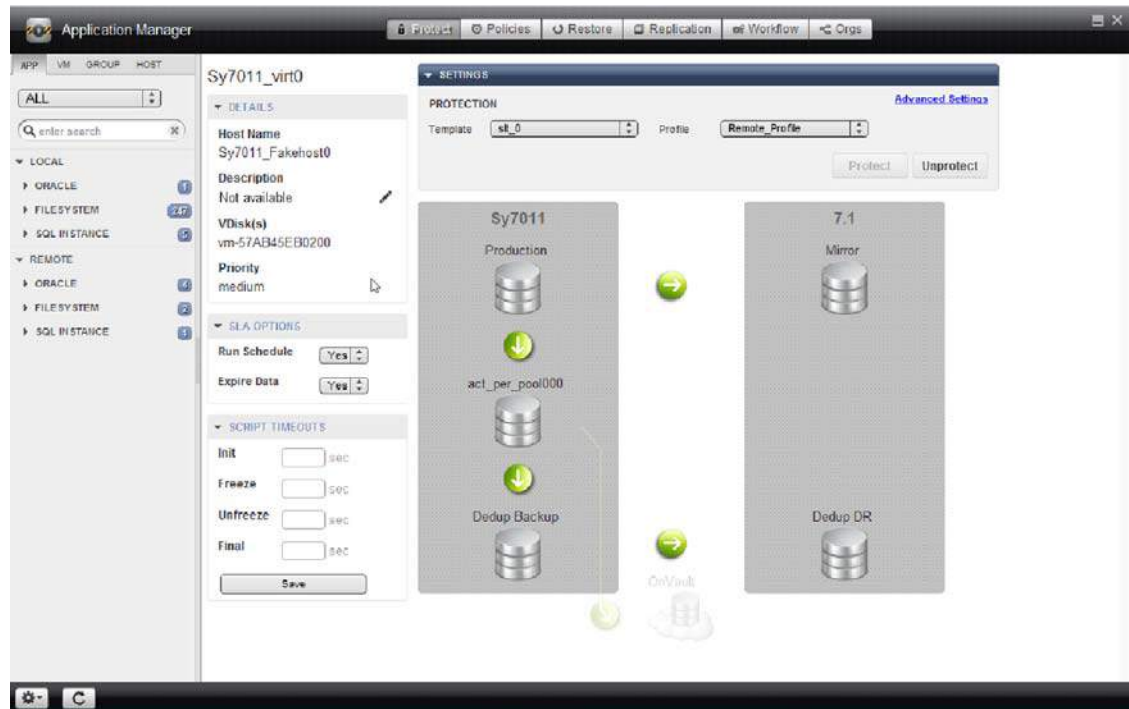
2. Select the local Actifio appliance from **Local** drop-down list. This is the appliance on which the profile is created.
3. When applicable, select the remote Actifio appliance from the **Remote** drop-down list.
4. Select the Snapshot pool from the **Disk Pool** drop-down list.
5. When applicable, select an Actifio OnVault Pool to which data will be sent. You can select this option only if this Actifio appliance has defined an OnVault Storage Pool. See **Configuring Resources and Settings With the Domain Manager** for details.
6. Click **Save**.

Applying Policy Templates and Resource Profiles

After defining at least one policy template and one resource profile, use the Application Manager service to apply them to an application or VM.

To apply a policy template and a resource profile to an application or a VM:

1. Open the **Application Manager**.



2. Select **APP** or **VM** from the Navigation Pane.
3. Select **Unprotected** from the Filter Pane and all unprotected (not captured) applications or VMs are displayed in the Navigation Pane.
4. Select the application or VM to capture.
5. Select a policy template from the **Template** drop-down list.
6. Select a Resource Profile from **Profile** drop-down list.
7. After applying a policy template and a resource profile, click **Advanced Settings** and the Advanced Settings dialog box is displayed. Advanced settings are specific to each application. If **Allow Override** is set **Yes** on the template assigned to this application, then the settings in the application's Advanced Settings dialog box will be applied.

Advanced Settings for specific applications are described in the **Virtualizing and Protecting Copy Data with the Application Manager** guide in the Actifio Documentation Library or the ActifioNOW customer portal.

4 Accessing Data

This chapter describes the various ways in which you can access your captured data:

[Mounts](#) on page 37

[Clones](#) on page 38

[LiveClones](#) on page 38

[Workflows](#) on page 38

[Restores](#) on page 38

For detailed, application-specific, step-by-step instructions on how to access data, refer to the following guides in the Actifio Documentation Library or the ActifioNOW customer portal.

- ***Accessing and Recovering Copy Data with the Application Manager***
- ***Replicating Data Using Actifio Appliances***
- ***Restoring Copy Data with the Application Manager***
- ***Creating Automated Workflows for SQL Server Databases*** and ***Creating Automated Workflows for Oracle Databases***

Mounts

The Actifio mount function provides instant access to data without moving data. There are two options for mounting data:

- **Application aware mounts** allow you to mount captured Microsoft® SQL and Oracle databases as virtual applications. This allows you to quickly bring a database on line without having to actually move the data and without having to manually configure a new instance of the database.

Application aware mounts are particularly useful in test and development environments where multiple copies of a database must be quickly brought on line.

Data presented as an application aware mount can be captured like any other application. Once the application aware mounted application data is captured, it too can be can be mounted as an application aware mount.

The capture, application mount, capture sequence can be repeated to any depth. By default, the sequence is restricted to five generations of the original database.

- The **standard mount** presents and makes application data available to a target server as a file system, not as an application. This is useful if an application is corrupt, lost, or if an application server is being replaced. In such cases you can mount an image and copy the application files from the mounted image to their original location on the application server.

Note: *Actifio Sky appliances can use a standard mount to access data that resides in an Actifio OnVault. Actifio CDS appliances **cannot** use a standard mount to access data that resides in an Actifio OnVault.*

Clones

Use the clone function to create an independent copy of a data set. The most common uses are: application development and testing, data audit for compliance, data warehousing, e-discovery, and user acceptance testing. Physical server or VM application-consistent data sets can be copied to a separate storage location anywhere in your environment. Like any other VM, a VM clone can be migrated to a new storage location.

LiveClones

The LiveClone is similar to the Clone function, however, unlike a Clone, a LiveClone can be updated on demand or according to a schedule. When an updated copy of the data is available, a LiveClone allows an independent copy of a data to be mounted. 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.

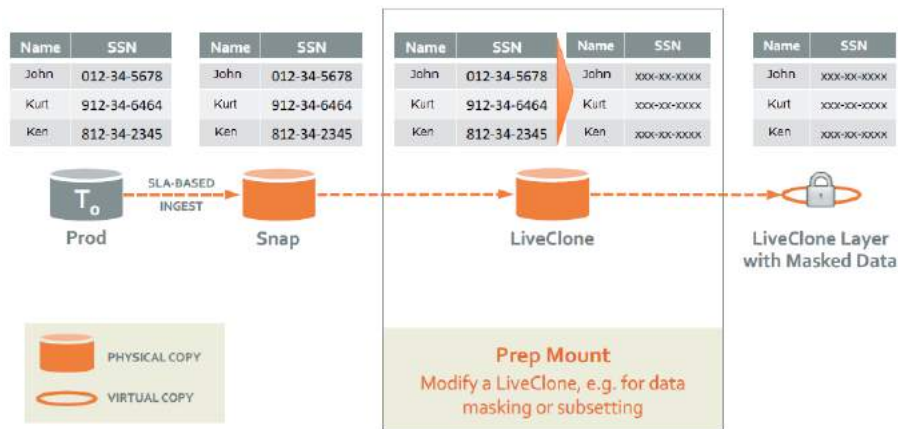
Workflows

Within the context of an Actifio appliance, a Workflow refers to the automation of access to copy data in a scheduled or prescribed fashion. While SLA Policy Templates govern the automated capture of production data and its management as copy data within the Virtual Data Pipeline, Workflows automate the access to this data.

Steps are defined within a Workflow to perform a series of tasks on a schedule or on demand. This includes creating and refreshing LiveClones, data masking, persistent mounts, and non-persistent processing mounts for tasks such as tape-out, database integrity checks, and ETL loads. Workflows are also used by administrators to provide simplified and secured self-service data access to end users such as database administrators and application developers.

For details on Workflows, see ***Creating Automated Workflows for SQL Server Databases*** and ***Creating Automated Workflows for Oracle Databases***.

The following image provides a high-level description of a Workflow that creates a LiveClone from production data, then scrubs the LiveClone of sensitive information, before mounting the LiveClone to a work environment.



Workflow With Scrubbed Data

Restores

The Restore function reverts the production data to a specified point in time. Restores and clones are the only data access operations that actually move data. Typically restore operations are performed to restore an application 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.

Glossary of Terms

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I-K](#) | [L](#) | [M](#) | [N-O](#) | [P-Q](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W-Z](#)

	Term	Definition
A	Actifio VDP™	Actifio Virtual Data Pipeline is the technology on which the company's unique copy data virtualization is built. See Virtual Data Pipeline™ .
	Actifio Connector	An Actifio Connector is a lightweight service that may be run on physical or virtual appliances. It discovers and captures individual applications virtual and physical machines and servers so they can be replicated.
	Actifio Desktop	The Actifio Desktop software controls the configuration and operation of a single physical or virtual Actifio appliance. See also Actifio Global Manager .
	Actifio Global Manager	The Actifio Global Manager (AGM) provides a web-based interface to manage multiple Actifio CDS and Actifio Sky appliances, including day-to-day copy data operations.
	Actifio Sky for AWS	Actifio Sky for AWS is a version of Actifio Sky software that creates and manages a virtual Actifio appliance in an Amazon AWS cloud space.
	Actifio Sky™	Actifio Sky is a cloud technology that creates and manages virtual Actifio appliances.
	AGM	See Actifio Global Manager .
	App or Application	An app or application is a data resource that can be discovered and protected by an Actifio appliance. Examples include Oracle or SQL databases, Exchange databases, network or local file systems or parts of file systems, virtual or physical machines, and so on.
	Appliance	An "appliance" is the generic term for either an Actifio physical or virtual server. Physical Actifio appliances are made up of two Actifio CDS or CDX nodes, a primary and a secondary and optional storage and fibre channel switches. Virtual Actifio appliances are referred to as Actifio Sky appliances.
	Application Manager	The Application Manager is software used to discover applications, application data, and virtual machines, and to apply protection templates and resource profiles them.
B	Baseboard System Identification	The base board system identifier (BBSID) is an arbitrary unique number that becomes part of a unique suffix for a node's World Wide Node Number and World Wide Port Number. These both must be unique within a fabric. The BBSID is also used by Actifio to generate a unique ID for Actifio SecureConnect access and to generate secure shell keys to access a system for troubleshooting.
	BBSID	Baseboard System Identification

	Term	Definition
	Actifio NAS Director	See NAS Director .
	Big Data Director	See NAS Director
C	CBT	Changed Block Tracking .
	CDS node	One of a pair of servers (a primary and a secondary) that makes up a physical appliance. A CDS appliance runs CDS software and VDP software and can manage fibre channel connected storage.
	CDX node	One of a pair of servers (both high-availability) that makes up a physical CDX appliance. a CDX appliance runs Sky and VDP software.
	Changed Block Tracking	Changed block tracking is the process of comparing the golden snapshot to incremental point in time snapshots in order to identify changed data that must be preserved.
	CLI	Command line interface.
	Clone	The clone function creates an independent copy of a data set. A virtual server or physical server data set can be copied from any application-consistent point in the system to a separate storage location anywhere in the environment.
	Clone VDisks	Clone VDisks, are the part of a Snapshot pool that contains full copies of an application's production data.
	Connector	See Actifio Connector .
	Consistency Group	A group of storage resources protected as a single entity by an Actifio appliance.
	Copy Data Virtualization	Copy Data Virtualization is the Actifio data management model — capture data and process it in a virtual data pipeline to create a single golden master copy that is incrementally updated according to a service level agreement (SLA) and is used to generate a virtual copy of any application data from any point in time for any authorized use.
D	DAR	See Dedup Async™ Replication (DAR) .
	Dedup	Deduplication is a storage technology and a process that reduces the amount of storage space consumed by data by removing redundant data.
	Dedup Async™ Replication (DAR)	Dedup Async Replication is a unique, proprietary Actifio technology that keeps a remote copy of production data always up-to-date and ready for data recovery. Also see Dedup Backup Replication .
	Dedup Backup Replication	An Actifio proprietary deduplication-aware replication protocol used for replication of captured images from one Actifio appliance to a second and optionally to a third for long-term storage. Also see Dedup Async™ Replication (DAR) .
	Domain Manager	The Domain Manager software controls the organizations and users that have access to an Actifio appliance, identify hosts that Actifio can protect, and manages the resources on which copy data resides.
E	Enumeration	Enumeration is the first phase Garbage Collection . Enumeration analyzes the catalog of deduplicated data to determine what data must be kept and what can be deleted. it is the first phase of Garbage Collection and is followed by Sweep .
F	Failback	Failback is the recovery process used when a primary system or data is restored to operation after a Failover . Also see Syncback .
	Failover	The process of using a secondary system, usually hardware, to replace a primary system that fails during operation. Also used to describe the data copied when a failover occurs. See Failback and Syncback .

	Term	Definition
	Fibre Channel	Fibre channel is a high-speed network technology commonly running at 2, 4, 8 or 16-gigabit per second that is used primarily to connect data storage devices.
	Filter Driver	The mechanism used by the Actifio Connector for Changed Block Tracking .
	Garbage Collection	Garbage collection is the two-phase process of reclaiming space in the dedup pool. Enumeration , the first phase, is the selection of data that can be deleted. Sweep is the removal of the unneeded data.
G	GC	See Garbage Collection .
H	Host	A server with managed or manageable applications.
	Hyper-V	Microsoft's virtual machine platform is a native hypervisor that can create virtual machines on x86-64 systems.
I-K	IGC	Incremental Garbage Collection was deprecated in Actifio CDS/Sky 7.0
	In-Band	In-band is a network architecture describing protected applications that are connected directly to Actifio storage. See also Out-of-Band and LAN-Free .
	iSCSI	The Internet Small Computer System Interface works on top of the Transport Control Protocol (TCP) and allows the SCSI command to be sent end-to-end over local-area networks (LANs), wide-area networks (WANs) or the Internet.
L	LAN-Free	LAN-free is a network architecture in which application data is protected using a shared, central storage device without sending the data over the local area network (LAN). See also In-Band and Out-of-Band .
	LiveClone	An independent clone of a captured image that consumes full storage resources and can be mounted to a host. It can be refreshed incrementally from another captured image, allowing very fast and efficient data refreshes for ETL and test & development purposes. A LiveClone can also be mounted for direct modification to support operations such as data masking.
M	Managed Data License (MDL)	Actifio's Copy Data Virtualization licensing. It is based on the amount of source data managed.
	Managed Disk	A SCSI Disk presented by a RAID controller and managed by the Actifio appliance. The Managed Disk is not visible to host systems on the SAN.
	Managed disk group (MDiskgrp or MDG)	A collection of Managed Disks that jointly contain all the data for a specified set of Virtual Disks.
	MDisk	These are disks presented to and managed by the Actifio solution
	Mount	The mount function is the most frequently used data access method, as it directly leverages the virtual copies of data stored on an Actifio appliance. By eliminating the data movement from the process, data sets of any size can be accessed instantly on any server.
	Multi-Hop replication	Replication, usually Dedup Backup Replication , is the process that replicates data from a "source" Actifio appliance to a "remote" Actifio appliance, and then from the remote appliance to a third Actifio appliance.
N-O	NAS Director	The Actifio NAS Director, formerly known as the Big Data Director (Actifio NAS Director) is a hardware and software product that creates an efficient platform to reliably capture, archive, replicate, and recover petabyte-sized unstructured datasets from large file systems.
	OnVault	The Actifio appliance vaults data to selected cloud storage according to a defined OnVault policy. Supported cloud storage platforms include Amazon S3, Google Nearline Storage, IBM Cloud Object Storage, and Microsoft Azure. Users manage and pay for their own cloud storage directly with the provider

	Term	Definition
	Out-of-Band	Out-of-Band is a network architecture describing protected applications that are housed on storage systems that are not connected directly to Actifio storage. See also In-Band and LAN-Free .
P-Q	Performance pool	The Snapshot pool .
	Policy	A policy defines when data will be captured, how long it will be retained, and where it will be replicated.
	Policy template	A collection of policies that, together, define when to perform a snapshot, when to perform dedup activity that creates an image, and how long to retain the image.
	PSRV	The platform service is a component of Actifio software that coordinates VDP services and functions.
R	RD	Resiliency Director is an optional product that works with CDS and Sky appliances to create and manage data that are part of disaster recovery services.
	Replication	Copying data to another Actifio appliance.
	Report Manager	The report manager is an optional stand alone software package that reports on data protection and recovery operations.
	Resource Profile	A resource profile specifies if, and which, Snapshot pool is used by Actifio and/or to which remote Actifio appliance data will be replicated. A resource profile is paired with policy templates to protect a specific application by the Application Manager .
	Restore	The restore function reverts the production data to look exactly as it did at the time of the data collection point. Typical use cases for restore would be to recover an entire server or application to a valid state after a massive data corruption or storage array failure.
	RM	Report Manager .
	RPO	A recovery point objective is the maximum period in which data might be lost from an IT service due to a major incident. See RTO .
	RTO	The recovery time objective is a period of time and a service level within which a business process must be restored after a disruption in order to avoid a break in business continuity. See RPO .
S	Service Level Agreement	An Actifio service level agreement is the linkage of a single policy template that defines when to perform actions, and a resource template that defines what storage resources are used by the actions.
	SideBand	See LAN-Free
	Sky	Actifio Sky™
	SLA	See Service Level Agreement .
	Smartcopy	Smartcopy is an Actifio-patented method for coping files and raw device data from protected applications onto Actifio staging disks and from mounted VDisks to host systems. Smartcopy does not have to be enabled.
	Snapshot	A snapshot is the process that captures and stores the state of a Snapshot VDisk as a Snapshot VDisk .
	Snapshot pool	The snapshot pool holds “golden copies” of application data for short-term retention. Data is instantly accessible and not deduplicated. Policies determine how long data is kept in the pool and when data is deduplicated and moved to another pool. The snapshot pool contains Staging VDisk , Snapshot VDisk , and Clone VDisks .

	Term	Definition
	Snapshot VDisk	A Snapshot VDisk is part of a Snapshot pool that preserves the state of Staging VDisk at specific points in time. Snapshots are retained according to a predefined protection policy.
	Staging VDisk	A Staging VDisk is part of a Snapshot pool that contains the Actifio golden copy of an application. It is retained for as long as an application is protected.
	StreamSnap	Direct replication of incremental snapshots from a local snapshot pool to a remote pool, supporting a much lower RPO compared with Dedup-Async replication. StreamSnap is used in high-quality, high-bandwidth IP networks. StreamSnap keeps a full virtual copy of the application on the remote side, available for immediate failover, test failover, or mount operations.
	Sweep	Sweep is the second phase Garbage Collection . It copies data that must be preserved from discrete areas of the dedup store into a new, contiguous area of the dedup store to both defragment storage media and to free areas for reuse. See Garbage Collection and Enumeration .
	Syncback	Syncback is the process that verifies data that has failed over to be valid before a Failback . Also see Failover .
	System Monitor	One of the services within the Actifio Desktop the monitors the process of jobs.
T	Tape-Out	Actifio Tape-Out is a legacy product for long-term storage and archival.
U	UDS	Universal data system, a component of Actifio's VDP architecture
V	VDisk	Also referred to as a <i>volume</i> . See Virtual disk .
	Virtual Data Pipeline™	The underlying architecture and processes of the Actifio data virtualization environment.
	Virtual disk	These are disks presented to applications by the Actifio solution that appear to host systems attached to the storage area network as a SCSI disk. Each VDisk is associated with one I/O group.
	VM	Virtual machine. Actifio supports both VMware and Hyper-V instances.
W-Z	Workflow	Actifio Workflows automate access to captured data. Workflows can run according to a schedule or on demand. Workflows present captured data as a LiveClone, a virtual application, or as just the application data.

