ACTIFIO FOR AWS

Modernize Your Backup and DR Using Actifio in AWS

actifio
Radically Simple

ACTIFIO FOR AWS

Modernize Your Backup and DR Using Actifio in AWS

What is Actifio?

Actifio virtualizes the data that's the lifeblood of business. We do it in more than 30 countries around the world. Our Virtual Data Pipeline™ technology enables businesses to protect, access, and move their data faster, more efficiently, and more simply. It decouples data from physical storage, much the same way a hypervisor decouples compute from physical servers. For enterprise-class backup modernization, self-serve instant data access, or service provider business transformation, Actifio is the first and only choice for radically simple copy data virtualization.

What can it do for you?

That depends who you are, and what problem you're trying to solve.

- Are you a **backup administrator** looking to modernize backups, manage workloads across private, hybrid, and public cloud, or eliminate your second site DR infrastructure by leveraging AWS on-demand?
- Are you an **enterprise architect** looking to lower your TCO by replacing multiple point tools for snapshots, backup, deduplication, long-term data retention, low RTO DR with a single platform in the cloud?
- Are you a **DBA** working to reduce backup windows and instantly recover multi-TB Databases?
- Are you an **IT Ops manager** looking to accelerate pre-production/staging testing in AWS?

If the answer is yes to any of the above, this solution brief will explain how Actifio can help you achieve all of the above by leveraging Amazon AWS.

Basic Enterprise Needs

All enterprise organizations have a range of data protection requirements:

1. **Local Backup**: Protect and retain data locally for a few days, weeks or months to achieve rapid local data restoration.
2. **Remote Vault**: Replicate data to a remote site to secure a second backup copy for long-term data retention (LTDR).
3. **Disaster Recovery**: In the event of a site disaster, recover applications at the remote site.

Large, global enterprises have been using Actifio software and hardware appliances for all of the above needs using internal infrastructure at both local and remote sites. However, an increasing number of enterprises are looking to replace their remote site with Amazon AWS. The next few sections will describe how to leverage Actifio in hybrid environments i.e. Actifio in customer premises and Actifio for AWS.

Actifio Architecture Overview

Actifio Virtual Data Pipeline (VDP) is a patented technology from Actifio that comes in two flavors: Actifio Sky software appliance and Actifio CDS physical appliance. Users can use Actifio Sky or CDS appliances on their premises and Actifio Sky in AWS.

FIGURE 1 **Architectural Overview**

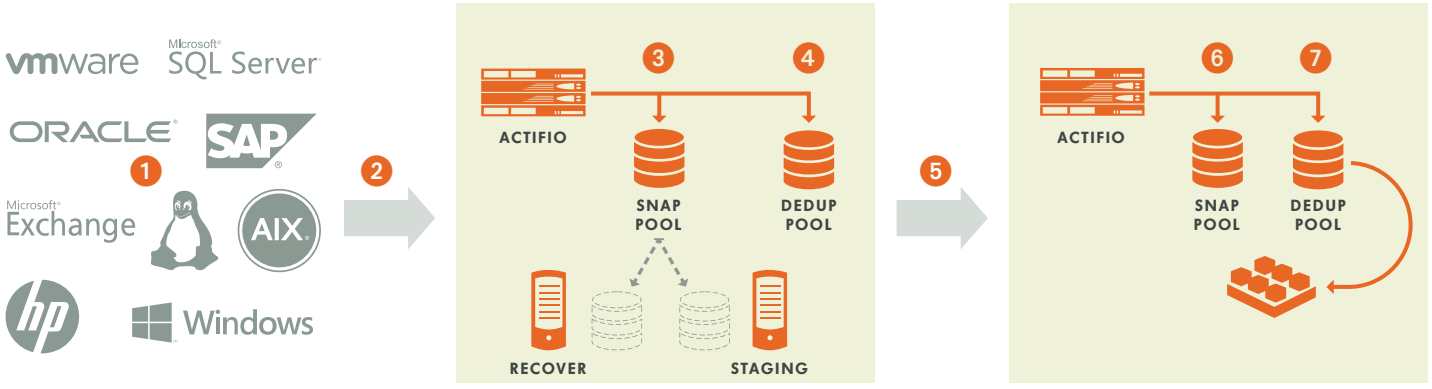

Actifio's high-level architectural data flow is shown in **Figure 1.**

1. Actifio has a wide platform support for many physical and hypervisor platforms, databases, and ERP applications.
2. Actifio Virtual Data Pipeline (VDP) does incremental forever backups at block level using native API integration such as VSS snapshots, Oracle RMAN APIs, VMware VADP APIs. After the first backup, Actifio ingests only changed blocks.
3. In native format, changed blocks are then stored in the "Snap Pool" for a short period. Actifio VDP then virtualizes Snap Pool data to enable instant recovery of the file system, DB, or VM.
4. Changed blocks are moved to the "Dedup Pool" and deduped across all applications for efficient long-term retention. This global deduplication reduces storage consumption and facilitates weeks or months of retention.
5. The deduped data is replicated efficiently over the WAN to another Actifio CDS or Actifio Sky appliance at remote DR site.
6. Data from multiple appliances replicating to a single target appliance is deduped globally and stored in the Dedup Pool. Users typically store data here for months or years, depending on their compliance and LTDR needs.
7. For selected applications that need low RTO DR, users can turn on "Dedup Async Replication" (DAR). Incremental backups for those applications will track and rehydrate blocks from the Dedup Pool to the Snap Pool at the DR site. Actifio virtualizes the DR-side Snap Pool data so that instant recovery is possible.

Actifio in AWS

Enterprise customers can deploy Actifio CDS or Actifio Sky on their premises and achieve local backup and local recoveries. Many enterprises also deploy Actifio CDS or Sky at their DR site to achieve Vaulting/LTDR. However, many of these enterprises are now looking at Amazon AWS as their remote site. Reasons include:

1. Eliminate all CAPEX on their own remote site
2. Elimination of operations time managing, upgrading and patching remote site infrastructure

The next section describes a few use cases and reference architecture of Actifio for AWS for each use case.

1. Vault To AWS
2. Low RTO Disaster Recovery (DR) In AWS
3. Pre-Production/Staging / UAT Testing in AWS

Vault to AWS

Figure 2 shows an architectural overview of how enterprise customers could vault their data securely to AWS.

1. Deduped data is replicated from a primary site to AWS over WAN. Actifio CDS/Sky encrypts data in flight over WAN. As of Sep 2015, all data coming IN to AWS is free.
2. Users would use their existing VPC (Virtual Private Cloud) or create a new one. For added security, an IPSEC VPN tunnel can also be created to encrypt data between firewall at primary site and the VPC in AWS.
3. Users can spin an appropriate EC2 instance and deploy Actifio Sky software in it. The choice of EC2 instance for Actifio Sky can depend on many factors such as amount of data protected, retention period, change rate etc.
4. Users can configure EBS (Elastic Block Storage) for Actifio Sky to configure the Dedup Pool. Users can choose EBS Magnetic volumes, or EBS General purpose (SSD) volumes, or Provisioned IOPS (SSD) volumes. When Actifio Sky is deployed in EC2, users can choose an option to encrypt data at rest. This will ensure that all data is encrypted on disk.

FIGURE 2 **Vault to AWS**

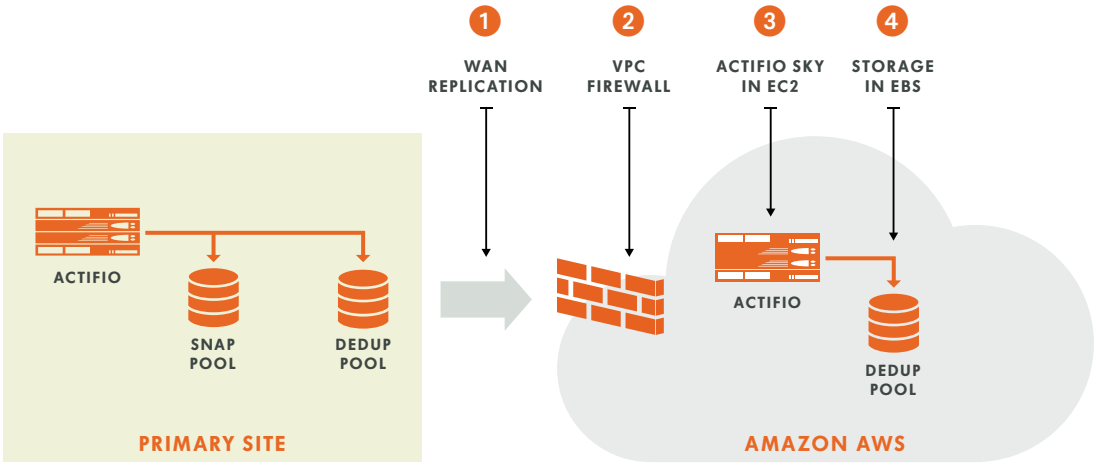

FIGURE 3 **Low RTO Disaster Recoveries in AWS**

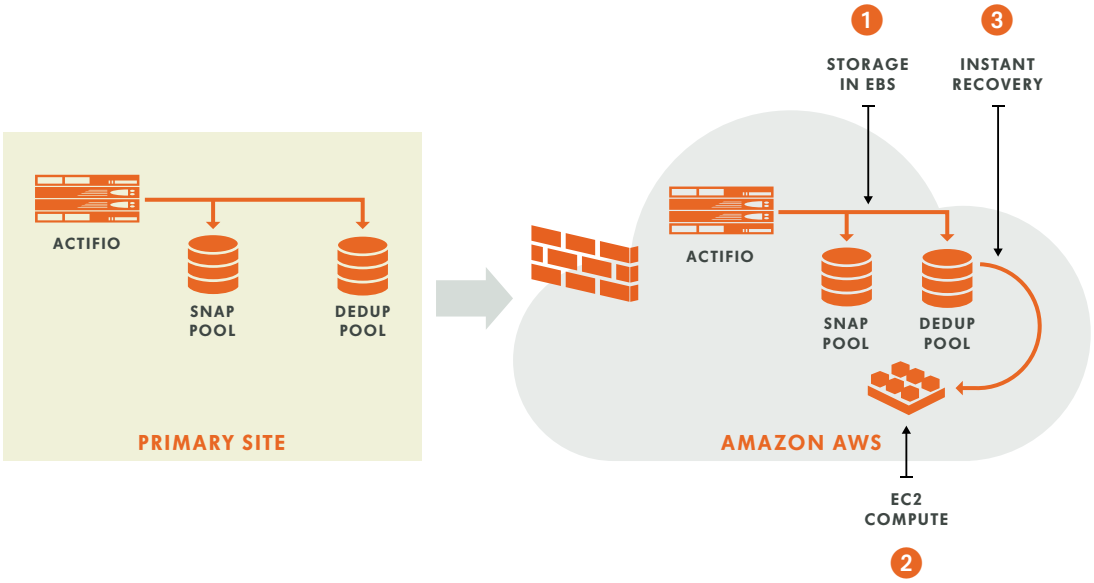

Actifio Sky in AWS stores deduped data in the Dedup Pool for long-term data retention. When users restore data from a very old point in time, most backup products must restore all data from remote dedup site over the WAN. However Actifio VDP first calculates the blocks that are needed to do restores. If, for example, it finds 60% of data blocks in the local Snap Pool, and 20% in the local Dedup Pool, Actifio will ingest only 20% from the remote AWS Dedup Pool. And these recalled blocks from remote site stay deduped and compressed over WAN. This ensures that the restores are not only fast, but also efficient so that the data OUT charges from AWS is as low as possible.

The following example illustrates approximate AWS sizing for an enterprise with 100 TB to be protected; an average change rate of 5% per day; 6 months retention for monthly backups. For these requirements, user would need the following:

1. Number of r3.4xlarge EC2 instances with Actifio Sky software = 4
2. Total TB storage needed in EBS (for Dedup Pool) to vault all data for 6 months = 200 TB
3. Bandwidth needed to replicate = 60 Mbps

Low RTO Disaster Recoveries (DR) in AWS

Actifio Vault in AWS, as explained above, already has the deduped data in AWS. Users just have to enable DAR (Dedup Async Replication) feature in Actifio CDS/Sky at their primary site to enable incremental forever rehydration of data from the Dedup Pool to the Snap Pool as shown (1) in the picture below. This ensures that data is already rehydrated and at the time of DR, no time is wasted in rehydrating data; thus reducing the RTO.

As a best practice we recommend users create EC2 templates in AWS based on production OS and application binaries. Thus, at the DR as shown in (2) you can spin up EC2 instances off these templates. This will ensure that the VMs have the right OS, service packs, patches and application binaries already installed.

The third step of disaster recovery uses Actifio to mount data volumes from Actifio Sky in AWS to the live EC2 VMs. This happens very fast because of Actifio data virtualization. Data volumes are mounted from Actifio Sky to the EC2 VMs over iSCSI. By the end of this step, users have their VMs up and running in AWS, thus completing the server recoveries in AWS.

From the example mentioned in Vaulting, assume that a user decides to have low RTO recoveries for 50 TB in AWS. This would mean additional storage in the Actifio Sky Snap Pool in AWS. For convenience the overall AWS sizing is shown below. The storage required for just low RTO is shown in **bold**.

1. Number of r3.4xlarge EC2 instances with Actifio Sky software = 4
2. Total TB storage needed in EBS (for dedup pool) to vault all data for 6 months = 200 TB
3. **Total TB storage needed in EBS (for snap pool) for low RTO DR = 65 TB**
4. Bandwidth need to replicate = 60 Mbps

More importantly for recoveries, users pay for compute on demand. Thus if they spin up, as an example, 500 EC2 instances for 48 hours, they pay for just those 48 hours. This eliminates a need to build a DR infrastructure in their remote data centers that is rarely used.

Pre-Production / Staging / UAT Testing in AWS

Enterprise users have to roll out new patches on existing production applications. These could be any of the typical OS upgrades, service pack updates, security vulnerability patches, application patches, new code drops etc. In all of these scenarios, enterprise IT typically tests the patches on a pre-production/staging environment, and only after successful completion of pre-production tests; they roll out patches on production.

In a typical test cycle, storage admins are requested to assign storage. System administrators are requested to provision appropriate physical or virtual systems with the right OS and application binaries. Backup admins are requested to restore data from production backups to pre-production environments. Then enterprise IT Ops deploy the latest patches and perform their functionality, security, and performance related functionalities.

Instead of having a standby environment, enterprise IT users can leverage AWS cloud for pre production/staging testing. Even without a disaster, enterprise IT users can spin up EC2 compute and mount Actifio data on that compute, as explained in the previous section "Low RTO DR in AWS". IT ops can use this environment as a pre-production/staging environment to deploy patches and execute their test cases.

This does not disrupt production, backups or replication. Once the IT ops are done with pre-production testing, they can spin down the environment in AWS.

This process of leveraging AWS and Actifio for pre-production testing has the following benefits:

- Reduce CAPEX by eliminating on premises pre-production infrastructure (compute, storage, network etc.)
- Reduce operations costs needed to manage, maintain the on premises pre-production infrastructure
- Reduce the pre-production testing time because of, a) speed with which EC2 VMs can be instantiated, and b) instant data mounts from Actifio Sky to these EC2 VMs

Summary

Actifio offers multiple benefits for many different enterprise users.

For Backup Administrators

1. SLA/policy-based data protection
2. Incremental forever backups, even for multi TB DBs
3. Instant recoveries, even for multi TB DBs
4. Low RTO DR in AWS
5. LTDR in AWS
6. Low bandwidth needs for replication
7. Wide platform support (physical, virtual, AWS, DBs, ERP Apps)
8. Central management and monitoring
9. Encryption in transit, Encryption at rest, Role-based access control, and Audit Trail

For DBAs

1. Reduce impact on multi TB production DBs because of incremental forever backups
2. Self-service, instant recovery of multi TB databases
3. Spin up multiple virtual copies of multi TB databases on premises or in AWS instantly using Actifio "app aware workflows"
4. Virtual database copies are writable and scalable
5. Execute reporting, analytics on virtual database copies on premises or in AWS on demand
6. Support for all major enterprise databases such as Oracle, Oracle RAC and ASM, MS SQL, SQL Clusters
7. Support for major ERP applications such as Oracle EBS, SAP, Microsoft Dynamics

For IT Managers/Directors/CIOs

1. **Lower TCO**: Eliminate multiple point tools such as storage snapshot, different backup products for physical, virtual, AWS; tapes, dedup appliances, DR infrastructure, storage and host based replication for low RTO, WAN accelerators
2. **Subscription**: Leverage pay-per-use Cloud Economics using Actifio for AWS

Try and Buy

A 30-day free trial of Actifio in AWS is available here.

For addition information, visit Actifio at actifio.com or email us at info@actifio.com or call us at **855-886-8997**