

CDS and Sky Tech Brief

Configuring Short RPO with Actifio StreamSnap and Dedup-Async Replication

Actifio recommends using Dedup-Async Replication (DAR) for RPO of 4 hours or more and using StreamSnap for RPO of 1 hour or more. This document describes best practices for using Actifio StreamSnap and DAR for offering short recovery point objective (RPO) SLAs. It is intended for system architects and managed service providers who need to understand how an Actifio appliance works and how to best size and configure it for consistent SLA compliance. This tech brief discusses:

[Best Practices and Considerations](#) on page 1

[StreamSnap Examples](#) on page 3

[Dedup-Async Replication Data Flow](#) on page 4

[RPO at the application level](#) on page 6

[Customer Examples](#) on page 8

Best Practices and Considerations

Measuring RPO

There are various ways to measure RPO, though all relate to the exposure for data loss. Since this document is about replication, we will only deal with “disaster recovery” scenarios, where the data has to be replicated to another site.


The “classic” and most stringent definition is “the longest period of time from which transactions could be lost in case recovery is needed”. For example, if you capture changes every hour and it takes an hour to copy the changes and replicate them to a remote site, then in the worst case you could lose 2 hours worth of transactions. This is because just before a replication job is finished the data that is available is almost 2 hours old. The RPO would be 2 hours.

A different measure of RPO is “the amount of time it takes to capture and replicate data”. So given the above example, in this measure the RPO would be defined as 1 hour. For simplicity, we will use this measure of RPO. If the above, stricter measure is needed, then you need to combine both the frequency and the time to capture and replicate, which could be as much as double the frequency of jobs.

RPO at the system level

When architecting and sizing a solution, the first step is to look at the entire system, in this case typically an Actifio appliance. The amount of data that needs to be handled is aggregated to find the total I/O performance, network bandwidth, and dedup engine throughput that are needed. Designing for a certain RPO is a classic flow problem where you need to ensure that the various components of the system along the data flow can handle the workload fast enough to meet the required RPO, that is, to move the data from start to finish within the allotted time.

StreamSnap Data Flow



StreamSnap Data Flow

1. The application is snapped or put in “hot backup” mode and incremental changed blocks are copied to the staging disk in the Actifio snapshot pool. The write operation into the snapshot pool storage, configured under Actifio management, is a copy-on-write (COW) operation, which generates 2 writes and a read for each changed block. As a general rule of thumb, to handle x amount of changed data you need 3x the raw storage I/O performance. For example, an array, storage controller, or RAID group that can handle 300MB/s of writes will sustain 100MB/s of changed application data.
2. The changed blocks are read from the snapshot pool and replicated over the network to the remote appliance.
This adds a read to the I/O load above, so the total demand from the underlying snapshot pool storage is 4x the amount of changed application data.

Note: Depending on the speed of the network, there might be some I/O caching that can reduce the load on the underlying spindles but it's best to use this as a conservative measure of the I/O requirements.

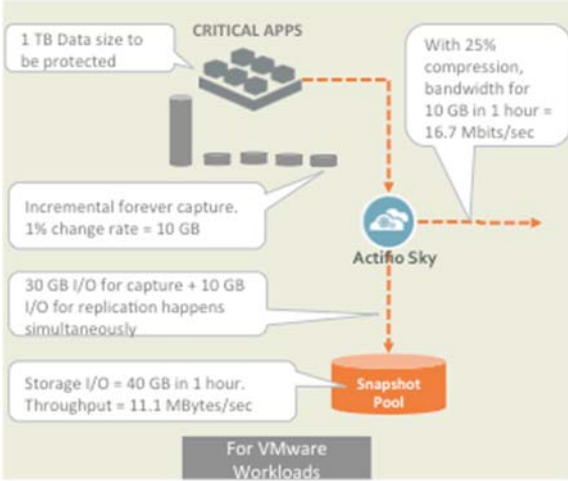
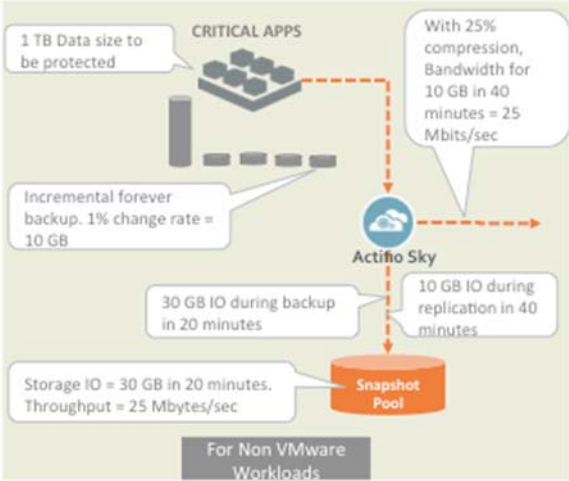
For most application types, this phase occurs after all the changes are copied to the snapshot pool, and therefore the time to replicate the data should be added to the data copy time from the previous phase, and both should fit within the required RPO. For VMware VMs, the replication starts in parallel with the copy to the snapshot pool and, if the network is fast enough, could proceed in parallel. An Actifio appliance can easily saturate a 10Gbps network link if the storage used is fast enough.

3. As they are replicated to the remote appliance, the changed blocks are written to the remote snapshot pool. The I/O requirements are similar to those for the local snapshot pool writes from phase 1.

StreamSnap Examples

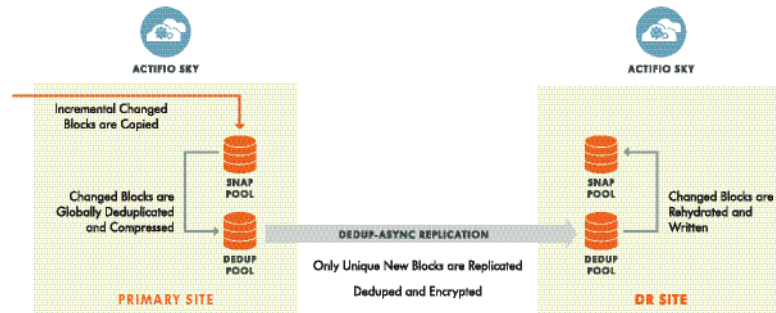
There are two distinct operations happening on the local appliance - incremental capture during phase 1 and replication of changed blocks during phase 2. How these are calculated depends upon whether you are working with VMware VMs or other applications.

Calculations for VMware and Non-VMware Applications

Calculations for VMware VM workloads	Calculations for non-VMware workloads
 <p>For VMware Workloads</p>	 <p>For Non VMware Workloads</p>
<p>For VMware VMs, replication starts in parallel while the incremental capture is happening. If the application being protected is 1 TB in size, the desired RPO is 1 hour, the average change rate per hour is 1%, and if we get 25% compression during replication, we calculate:</p> <ul style="list-style-type: none"> • The average amount of data changed in an incremental capture is 10 GB (1% of 1 TB). • The underlying storage assigned to Actifio snap pool should be able to handle a storage I/O of 4x of data changes, so $4 \times 10 \text{ GB} = 40 \text{ GB}$ in 1 hour, equal to storage throughput of 11.1 MegaBytes/second. • After compression of 25%, 7.5 GB needs to be replicated in 1 hour, i.e. network bandwidth needed is 16.7 Megabits/second. • The storage assigned to the Remote Actifio snap pool must keep up with the incoming replication data and writes. Since these are COW writes, the required I/O is 3x the amount of changed data (10 GB uncompressed), so the I/O needed is $3 \times 10 \text{ GB} = 30 \text{ GB}$ in 1 hour, equivalent to storage throughput of 8.3 MegaBytes/second. 	<p>For all other application types, replication starts after the incremental backup is finished. In this example, the storage used for Actifio snap pool should be able to handle 30 GB of I/O during the incremental capture phase and 10 GB of I/O during the replication phase. Since bandwidth is expensive, we optimize for bandwidth, designing the system for 20 minutes of capture and 40 minutes of replication. We calculate:</p> <ul style="list-style-type: none"> • The storage throughput requirements are 30 GB in 20 minutes, which is equivalent to 25 MegaBytes/second. • The network bandwidth requirement with compression is 7.5 GB in 40 minutes. i.e. 25 Megabits/second. • The storage assigned to the Remote Actifio snap pool needs to keep up with the incoming replication data and writes. As explained, since these are COW writes, the required I/O is 3x the amount of changed data (10 GB uncompressed). So the I/O needed is $3 \times 10 \text{ GB} = 30 \text{ GB}$ in 40 minutes (the period for replication), which is equivalent to storage throughput of 12.5 MegaBytes/second.

The key variables are the average change rate per period, storage performance, and network bandwidth. Also important is knowing that change rates fluctuate over the course of the day. Designing storage and bandwidth for a peak change rate will increase adherence to SLAs as well as costs. For example, a 1 TB VM with a peak change rate of 5% between 2pm to 3 pm. The amount of changed data is 50 GB (instead of 10 GB) and this needs to be replicated in 1 hour between 3 pm to 4 pm. The bandwidth to handle peak change rate (assuming 25% compression) for 1 hour RPO is then 83 Megabits/second, five times higher than the bandwidth designed for average change rate. Architecture for an average change rate will save on costs but you must expect to miss the RPOs for some hours.

Dedup-Async Replication Data Flow



DAR Data Flow

1. The application is snapped or put in “hot backup” mode and incremental changed blocks are copied to the staging disk in the snapshot pool. This is a copy-on-write (COW) operation, which generates 2 writes and a read for each changed block. As a general rule of thumb, to handle x amount of changed data you need 3x the raw storage I/O performance. For example, an array that can handle 300MB/s of writes will sustain 100MB/s of changed application data.
2. The changed blocks are processed by the dedup engine for global deduplication and compression, and written to the local dedup pool. This adds a read to the I/O load above, so the total demand from the underlying snapshot pool storage is 4x the amount of changed application data.

Global deduplication and compression done by the dedup engine can result in significant reduction of required bandwidth but adds processing time. Therefore, the throughput of the dedup engine is a key determinant in the amount of data that can be handled and in the achievable RPO. The dedup engine performance is largely driven by the performance of the dedup pool storage - general guidelines are given in the table below.

3. The changed blocks are also compared with the remote appliance's dedup engine to identify those blocks that need to be replicated. This is done very efficiently using block hashes and other optimizations that ensure that only unique blocks of data that have never been replicated are actually sent over the wire.
4. On the remote appliance the dedup engine writes the new data blocks to the dedup pool as well as rehydrates and writes all the changed blocks to the remote snapshot pool. The dedup engine throughput and workload are similar to those on the local appliance (phase 2) and the snapshot pool I/O requirements are similar to those for the local snapshot pool writes (phase 1).

General guidelines for dedup engine throughput and the storage required to support it

Appliance:	CDS	Sky 50TB	Sky 30TB	Sky <30TB
<p>Tier 1 - SSD</p> <p>Required IOPS (of 64KB): 15,000 random reads and 15,000 sequential reads with 50 threads; 15,000 random writes and 15,000 sequential writes with 100 threads</p> <p>Reference system: Pure Storage AFS</p>	12 TB/day	9 TB/day	6 TB/day	N/A
<p>Tier 2 - 10K/15K RPM SAS</p> <p>Required IOPS (of 64KB): 4,000 random reads and 15,000 sequential reads with 50 threads; 2,000 random writes and 10,000 sequential writes with 100 threads</p> <p>Reference system: NetApp E5600 30x800GB 10,000 rpm SAS drives, RAID-6 DDP</p>	8 TB/day	6 TB/day	4 TB/day	3 TB/day ^a
<p>Tier 3 - 7,200 RPM SAS</p> <p>Required IOPS (of 64KB): 2,200 random reads and 4,000 sequential reads with 50 threads; 380 random writes and 2,800 sequential writes with 100 threads</p> <p>Reference system: IBM StorWize v3700 12x4TB 7200 rpm SAS drives, RAID 6</p>	4 TB/day	3 TB/day	2 TB/day	1.5 TB/day ^b

- All throughput numbers in this table include the use of SSD for dedup index acceleration. This is required for Sky 50TB and Sky 30TB, and optional for Sky models smaller than 30TB. If no SSD is used for dedup acceleration for these smaller models then throughput numbers will be lower than those listed here.
- The above throughput numbers are for system-wide performance and assume that work can be sufficiently parallelized to leverage all system resources. In particular, each volume of each application is handled as a separate thread. For CDS and Sky 50, there are 7 cores that are available for concurrent dedup work (so 7 volumes can be deduplicated in parallel), while fewer cores are available in smaller Sky models.

RPO at the application level

While designing and sizing for meeting a certain RPO, care must be taken to look at individual applications, especially the larger ones. The system might be sized properly to handle all the data in aggregate, but if some applications are much larger than others or have a much higher change rate then the RPO might not be met for those in particular. It is recommended to take the application that will generate the most change and ensure that its data can flow through the system within the desired time frame.

As mentioned above, the work unit within the system is a volume. If an application has multiple volumes then you should take the largest one and calculate the data flow for it. For example, for large Oracle databases the system will create multiple staging disks (typically 1TB or 2TB each) and combine them together as one large volume when presented to the production host. These staging disks can be processed in parallel (in sub-jobs), to leverage the full capabilities of the system resources. However, take into account that, depending on operation windows of different templates and the applications associated with them, other applications might be processed in parallel, limiting the resources that are allocated to the largest application (or that with the most change).

Calculating the time to write the changes to the local and remote snapshot pools and the time to replicate the data over the network is straightforward based on storage IOPS and network bandwidth. For DAR, keep in mind that a CDS appliance has 7 cores that are dedicated to dedup and so for a given sub-job handling one volume the throughput would be at most 1/7th of the total system throughput. For example, if using 7,200 RPM SATA drives, your throughput for one volume will be 4 TB/day divided by 7, which is 570 GB per day.

Variables that impact SLA compliance

There are some common variables that can significantly impact your SLA compliance as far as RPO:

Impacts:

- **Network congestion** - other activity on the network can obviously slow your data flow. Make sure that you can consistently provide the bandwidth that is used for RPO sizing.
- **Other running jobs** - if you intend to perform frequent mounting of virtual clones (e.g., for test data management) or expect to frequently have to restore data from deduplicated images, take into account the I/O impact of those jobs. Recovery jobs such as mount or restore typically have the highest priority in the system and will take away resources from ongoing capture jobs.
- **Higher than normal change rate for an application** - sometimes users perform actions that can drive dramatically high change rates. Re-indexing a database or performing an ETL operation into a database, for example, can cause massive changes at the block level. Sometimes users move applications (e.g., between data centers) and so a full new ingest of an application might be needed. You need to decide up front how much of such change you want to be able to handle while still meeting the desired RPO. Service providers often exclude SLA compliance if the customer generates changes that are more than some agreed-upon amount.
- **Loss of changed block tracking** - Sometimes, changed block tracking is lost (e.g., when an ESX server crashes). While Actifio has mechanisms to deal with that and avoid a full new ingest by scanning the production data and rebuilding the changes, that takes longer than an incremental capture. You need to decide if you want to account for that in the system sizing.
- **Distribution of change** - Finally, change is usually not distributed evenly. While the system might be sized to meet the RPO for the average change rate, and in aggregate that works fine, this could easily be violated, especially for larger applications. You might have days with more activity (on a weekly or monthly, or even annual basis), and even within a day change is often distributed unevenly. If you only size the system to the average change, take into account that you will get violations and account for that in how the SLA is defined and measured.

General Considerations and Suggestions

- In large deployments, dedicate an appliance for short RPO applications - Most customers protect only a subset of their applications with short RPO. When designing large deployments with multiple appliances or when accounting for very large applications with high change rate (TBs of change per day), consider dedicating appliances for these applications. This will isolate them from impacts of other applications that might be handled less frequently and could create "lumpy" behavior where the system gets busy with a large amount of data from an application that does not require the short RPO handling. If overall capacity under management is not very big, the use of a Sky appliance for this purpose could be an option. Isolating short RPO workloads, while certainly not a requirement, can often simplify deployment design and ongoing management.
- Prioritize short RPO applications - If there is a need to mix short and long-RPO applications in the same appliance, typically for small overall capacities, make sure to designate SLA template policies for short RPO with High priority, so that the system gives preference to jobs for short-RPO applications.
- Design to peak - Avoid sizing appliances at or very close to their processing capabilities so that temporary peaks do not put them "over the edge". Monitor system load, especially dedup engine slot utilization to ensure it does not go beyond 60% on average. In general, in order to ensure meeting SLAs you should "design to the peak" and not to the average. This applies to network bandwidth as well.
- Account for applications with different volume sizes - Take into account that for a given application some volumes might be much larger than others. For example, suppose your largest VM has 2 volumes - a 40 GB boot volume and a 2 TB data volume. A common mistake is to look at the total capacity (2.04 TB) and divide that by the number of total volumes (2 in this case). Doing that will give you an average volume size of ~1 TB (remember that the "work unit" for the system is a volume). However, if you base your RPO calculations on this average volume size, making sure you can handle 1TB of data in the RPO period, you'll be about 50% wrong, since the data flow will be gated by the 2 TB volume.
- Optimize resources for StreamSnap - If you design a system for snapshots and StreamSnap use only, without dedup, then there are further optimizations that could be done for Sky appliances. The standard Sky hardware specification allocates considerable resources to the dedup engine (CPU and memory). In cases where dedup is not used, you can provision a "small-dedup" Sky such as the Sky 1TB model, but allocate to it the same amount of CPU and memory that is spec'ed for a Sky 50TB (12 cores instead of 1 and 72GB of memory instead of 6GB). The additional resources will be used to speed up snapshot and StreamSnap operations. On the flip side, if you are trying to design a system that will have to handle high loads of StreamSnap replication but also deduplication (local or replicated remotely), consider provisioning additional cores and especially memory, beyond the minimum in the spec, to allow for further optimized performance and more adherence to strict RPO requirements.
- Configure storage properly - For storage configuration, each pool should be on separate RAID groups, and each appliance should be on a dedicated storage controller (this might not be necessary on higher-end storage arrays). Auto-tiering, such as with SSD, can be used for snapshot pools and SSD can be used as a buffer to help absorb irregular change rates. However, auto-tiering storage should not be used with the dedup pool due to the random write nature of the dedup activity, which will defeat auto-tiering mechanisms and generate unnecessary overhead and even thrashing. Instead, SSDs are used for accelerating dedup index lookups by allocating them directly to the dedup engine (this is required for CDS and for Sky 30 and Sky 50, and is optional for smaller Sky models).

Customer Examples

This section provides examples of some customer scenarios where DAR and StreamSnap are used to deliver short RPOs.

Large insurance fund - StreamSnap with frequency of 1 hour

- Actifio appliance protecting 10 applications, total MDL 50TB.
- Most applications are 3 to 8 TB in size, one is 20 TB.
- All but one are protected with a 1-hour StreamSnap frequency policy.
- Change rate is 1%, or about 0.5 TB per day.
- In a sample month, 98.6% of StreamSnap replication jobs completed within 60 minutes, with an average job run time of 11.3 minutes.
- Within that period, 98.1% of the StreamSnap replication jobs started within 80 minutes of the start of the previous job.

Medium size bank - DAR with frequency of 2 hours

- Actifio appliance protecting 200 applications, total MDL over 20TB.
- Of that, there are 13 applications with a 2-hour DAR frequency policy.
- Average size of these applications is 133GB, with 2.8% change rate.
- Within a sample week, ran 1,132 jobs with an average job length of 12.6 minutes.
- 99.2% of jobs finished within 4-hour SLA (meeting a “strict” definition of 4 hour RPO, meaning max loss of 4-hours worth of data).
- Excluding one problematic application, 99.9% of jobs finished within 4-hour SLA, with only one violation, by 8 minutes.

Large private equity firm - DAR with frequency of 4 hours

- Actifio appliance protecting 118 applications, mostly VMs, with total MDL over 11TB.
- Of that, there are 22 applications with 4-hour DAR frequency policy.
- Average size of these applications is 73GB (largest is 238GB), with 0.7% change rate.
- Within a sample week, ran 941 jobs with an average job length of 10.2 minutes.
- 100% of jobs finished within 8-hour SLA.