IBM Spectrum Virtualize
Data Reduction
Best Practices

Contents

Data Reduction **Techniques**

Thin Provisioning

Capacity is allocated on demand when the data is first written.

Compression and De-duplication make use of thin provisioning.

Compression

Data is compressed before being written to storage*

*Some other vendors don't run these operations in real-time. i.e. the original data is first written uncompressed/non-deduplicated and a background process reduces later.

Deduplication

Duplicates of data are detected and replaced with references to the first copy*

Spectrum Virtualize
Data Reduction Technology

Data Reduction with **Data Reduction Pools**

Data Reduction Pools (DRP) deliver **thin provisioning, compression** *and* **deduplication** across all volumes in a pool

Requires best practices be followed because :

- Garbage discounted at a DRP level may still exist in a physical storage level
  (Garbage is data that is no longer needed but has not yet been deleted)
- Using fully allocated volumes with a compressing back end may lead to a distorted view of physical storage remaining.
  (DRP sees a fully allocated volume, compressed back end sees a thin, compressed volume)

If best practices are not followed :

- Running out of storage will take volumes offline, making it difficult or time consuming to recover
- Under-estimating the amount of storage remaining negates the cost benefits of using data reduction and makes long term planning challenging

As with any data reduction technology, you should *always track usage of underlying storage*

Data Reduction with **FlashCore Modules**

Members of the **FlashSystem family(*)** contain compressing physical storage (ie FlashCore Modules, or FCMs)

FCMs compress data at line speed with *no performance impact to the workload*

DRP is fully supported within these products, even when using FCMs

If virtualizing external storage using these or any other Spectrum Virtualize products, then the best practices still apply!

When using FCM based products :

The user should **ALWAYS** monitor physical space

The GUI will ensure the user creates configurations that are sensible and measurable

And the storage controller uses knowledge of both the logical space and physical space
(This is an advantage of Spectrum Virtualize being tightly integrated with the storage controller where more information about the underlying storage is known)

(* FlashSystem family members include FlashSystem 9200, FlashSystem 7200, FlashSystem 5100, and the previous members **FlashSystem 9100, Storwize V7000 Gen 3** and **Storwize V5100)**

Considerations When Using Data Reduction at 2 levels

If you create a solution where data reduction technologies are applied at both the storage and the virtualization appliance levels, then here are the rules you should follow :

ALL DRP volumes should run with **compression switched on**
(DRP uses Intel Compression Acceleration hardware & will not adversely affect general DRP performance)

Physical storage should be allocated 1:1
This means the usable capacity in the storage must be equal to the effective capacity provisioned to SVC.
Review documentation for how to do this for your storage system.

Fully allocated volumes should be in their **own pool**

If you want to use DRP with an existing over-allocated backend, you need to reclaim storage and configure it according to the best practices

Data Reduction
Design Patterns

DRP above simple RAID

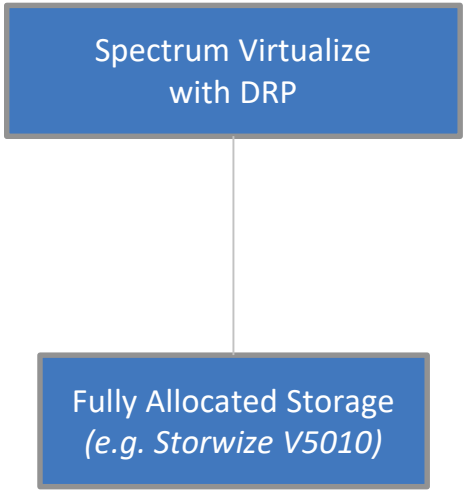

Recommended!

Use DRP at top level to plan for de-duplication and snapshot optimisations.

DRP at top level provides best application capacity reporting (volume written capacity).

Always use compression in DRP to get the most benefits from the technology

Fully-allocated above single-tier data reducing backend

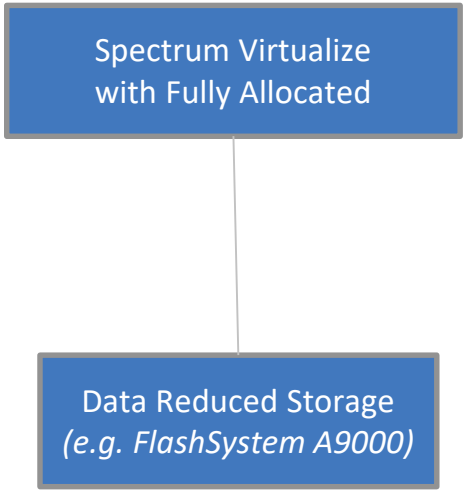

Recommended with appropriate precautions!

Need to track physical capacity use carefully to avoid out-of-space.

Spectrum Virtualize can report physical use but does not manage to avoid out-of-space.

No visibility of each applications data usage at the Spectrum Virtualize layer.

If actual out-of-space happens there **is limited ability to recover with unmap**. Also consider creating sacrificial emergency space volume.

DRP above data reducing backend

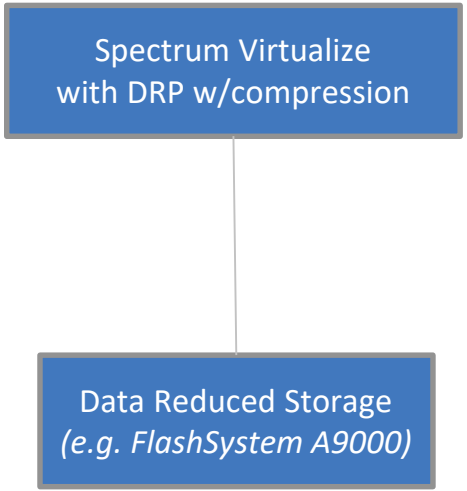

Recommended with appropriate precautions!

Assume **1:1** (effective:usable) compression in backend storage – **do not overcommit!**

Review product documentation for how to do this

Small extra savings will be realised from compressing meta-data.

Monitor capacity and out-of-space alerts from the backend storage system.

Using DRP with over-allocated back end could lead to the DRP garbage causing out-of-space

Think : For existing systems, do you need to move to DRP to get the benefits of dedup, or is line-speed hardware compression the right performance/capacity balance?

DRP + Fully-Allocated above data reducing backend

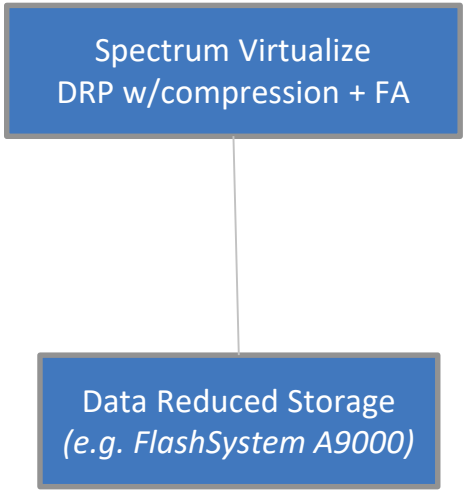

Avoid!!

Very difficult to understand used capacity when combining DRP and FA and a compression backend

Temptation is to try to exploit capacity savings which might ***overcommit backend***

If a mix of DRP and FA is required, then use separate pools; Previous design patterns can be used, but must not combined in the same pool

(Or use with *extreme* caution and with a ~<20% mix of fully allocated)

Fully-allocated above multi-tier data reducing backends

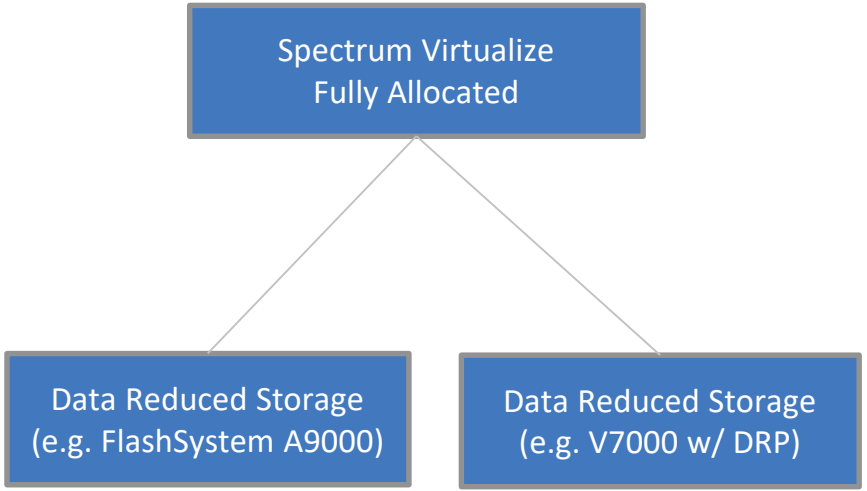

Use with great care!

Easy Tier is unaware of *physical* capacity in tiers of a hybrid pool.

Easy Tier will tend to fill the top tier with hottest data.

*Changes in compressibility of data between the tiers can overcommit the storage leading to **out-of-space.***

Think: Is there risk my hot data is incompressible?

DRP above DRP

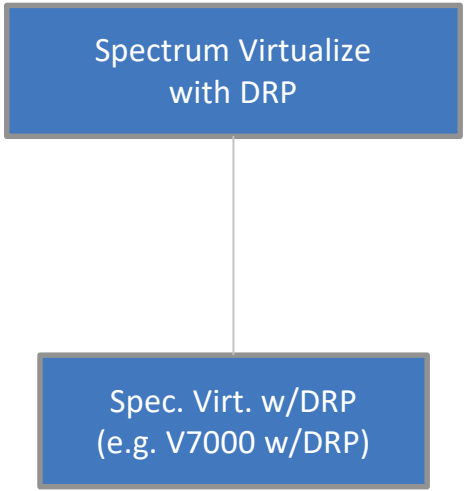

Avoid!!

Creates two levels of IO amplification on meta-data.

Introduces an additional layer of capacity management (two DRP, one physical)

There is no performance, capacity or usability benefit in using DRP on top of DRP

Data Reduction Usage
Further Considerations

Summary - **Planning**

Flexibility Is Great, But Choose The Right Technology

Use **Storage Modeller** to model the performance with both FCM and DRP data reduction
(IBM or your Business Partner can help with this)

Use **Compresstimator** to identify expected space savings. **Data Reduction Estimation Tool (DRET)** can identify deduplication workloads.

Use **FlashCore Modules** with inline hardware compression for best performance

Use **Data Reduction Pools** if you can really benefit from deduplication capacity savings

Choose The Right Design Pattern

Don’t over-complicate

If virtualizing a compressing storage system use a recommended design

Other designs make it impossible to understand your real capacity consumption and you will run out of space!

Summary - **System Sizing**

Right Sizing The Cache

Appropriately sizing the cache can improve the operation of the system

- **256GB** is a good starting point to allow a reasonable size non-DRP working set in cache
- Another **128GB** is recommended if you're making heavy use of copy services
- If consolidating workloads, consider how much cache you currently have available
- **Maximise** the cache to extract the best performance out of DRP (lower latencies, avoid response time spikes) and optimize DRP metadata hits

Right Sizing Your Capacity

As a rule, try to avoid running your system past 85% full

Workloads can be unpredictable and a large demand for capacity could quickly fill up system capacity. 15% buffer provides opportunity to handle such situations

As volumes fill, garbage collection work could trigger write amplification placing additional write workload on the system, slowing down application response times

Monitor your free space and plan ahead to avoid running garbage collection too frequently and fully realise the performance required by your workloads

Summary - **Deduplication Workloads**

Good Deduplication Workloads

A workload with a lot of identical data across one or more volumes

The more duplicates there are, the greater the savings realised in terms of physical space used

As mentioned, **data reduction costs IOPS**, so balance the capacity savings against the performance requirements

A great example of a deduplication friendly workloads is VDI

Successful VDI Deployment

VDI tends to have a large write contingent of around 30/70

As well as having a lot of duplication, this kind of workload benefits from having more cache

To fully realise the benefits of the system, configure the VDI deployment across multiple VMFS datastores so as to increase parallelism. At least 16 if possible.

Summary - **Monitoring**

Monitor Capacity

Configure alerts – get warned EARLY when capacity starts to fill up

- Allow enough time to delete or migrate data or augment existing capacity

Running out of physical capacity will take volumes offline

Monitor your reclaimable capacity and free space to do longer term planning

Monitor System Performance

Plan for high bandwidth demanding workloads (such as migration in) consuming excessive system resource

- Follow Spectrum Virtualize best practices to create a balanced system
- Consider additional resource usage when using DRP
- Use QoS to throttle heavy workloads and protect other workloads running on the system
  (Throttling may elongate time taken to run a workload)

Space savings need IOPS to be realized

Further Reading

White Paper - Best Practices Details for Managing Physical Space on FlashSystems 900-AE3

https://www-01.ibm.com/support/docview.wss?uid=ibm10735459