Driving IT Storage Effectiveness with Real-time Compression

Leveraging Data Footprint Reduction (DFR) and Storage Services Innovation

By Greg Schulz

Founder and Senior Advisory Consultant, the Server and StorageIO Group
Author "The Green and Virtual Data Center" (CRC) and
"Resilient Storage Networks: Designing Flexible Scalable Data Infrastructures" (Elsevier)

November 24, 2010

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com

Introduction

Storage services (Figure 1) continue to evolve from RAID, to virtualization, thin or dynamic provisioning, local and remote data replication, space saving snapshots, access (block, file, object and cloud) along with tiering and data footprint reduction (DFR). A focus on storage services is enabling efficient and effective storage optimization, boosting both performance and productivity across all application and tiers of storage.

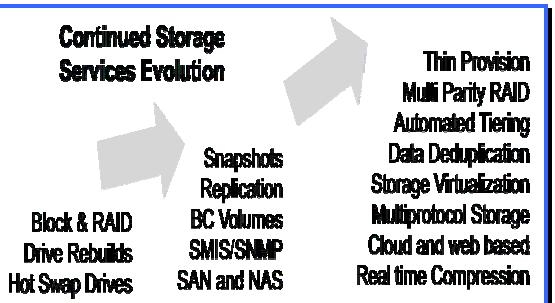

Figure 1: Storage Services Evolution

Background and Challenges

There is no such thing as a data or information recession! Most organizations are experiencing strong data storage growth and an expanding data footprint due to the need to store more data for longer periods of time. The result is that data storage infrastructure and associated management tools, best practices, budgets and resources are being strained. Consequently, IT organizations are looking for opportunities to stretch further available resources (people, time, budgets, physical floor space, server, networking, storage along with power and cooling).

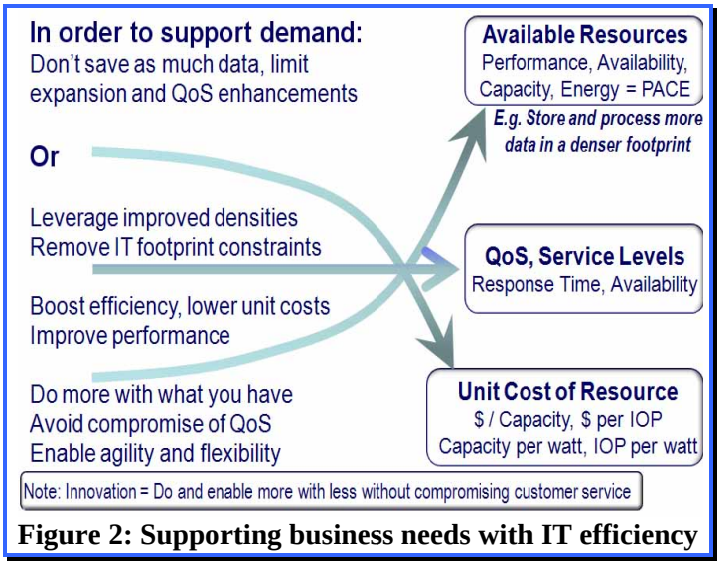

Figure 2: Supporting business needs with IT efficiency

More data is being generated, processed, moved and stored in multiple locations for longer periods of time, and this trend shows no sign of decreasing. Figure 2 shows a business common objective of a) reducing per unit cost of IT resources (servers, storage, networking, hardware, and software) while b) supporting increasing business demands. This needs to be done at the same time as stretching available budgets further (doing more in an efficient effective manner) and supporting demand without negatively impacting IT services delivery.

This means that more data needs to be managed, protected, preserved and served in the same or smaller footprint constraints (e.g. power, cooling, floor space, management tools, budget, data protection windows, skilled IT staffing) while maintaining or enhancing Quality of Service (QoS) and meeting Service Level Objectives (SLOs). Simply cutting costs is not sufficient for most environments as that could negatively impact quality of service (response time, responsiveness, availability, data protection, etc.) or impede business productivity needs. As noted at the bottom of Figure 2, innovation - doing more with what you have while reducing per unit costs to support increased demand without compromising QoS or SLOs for business or customer enablement - is needed!

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com

Understanding your data footprint and its impact

If it was not becoming increasingly necessary to process and store more data for longer periods of time in different locations, there would be little need for data storage, networking, IT clouds or virtualization. Likewise, if there was not the ever increasing dependence on information being accessible when and where needed there would not be the need for business continuance (BC), Disaster Recovery (DR) and backup/restore along with archiving.

Data footprint impact is the total data storage needed to support various business application and information needs. Your data footprint may be larger than how much actual data storage you have, as in the example shown in Figure 3. This example is an organization that has 20TBytes (add as many zeros to fit your environment's need) of storage space allocated and being used for databases, email, home directories, shared documents, engineering documents, financial and other data in different formats (structured and unstructured) as well as varying access patterns.

Figure 3: Expanding data footprint impact Source: The Green and Virtual Data Center (CRC)

Common IT costs associated with supporting an increased data footprint include:

- Data storage hardware and management software tools acquisition
- Associated networking or IO connectivity hardware, software and services
- Recurring maintenance and software renewal fees
- Facilities fees for floor space, power and cooling along with IT staffing
- Physical and logical security for data and IT resources
- Data protection for HA, BC or DR including backup, replication and archiving

The larger the data footprint the more data storage capacity and performance bandwidth needed. How the data is being managed, protected and facilitated (powered, cooled, and situated in a rack or cabinet on a floor somewhere) also increases the demand for capacity and associated software licenses. For example, in Figure 3, storage capacity is needed for the actual data as well as for data protection using RAID, replication, snapshots and alternate copies including disk to disk (D2D) backups. In addition, there can also be overhead in terms of storage capacity for applying virtualization or abstraction to gain additional feature functionality from some storage systems along with reserve space for snapshots or other background tasks.

Driving IT effectiveness with Data Footprint Reduction (DFR) innovation

Traditionally, if costs are reduced, one or more of the curves shown in Figure 2 can be negatively impacted. Thus the common tactic of moving problems around results in them resurfacing elsewhere possibly magnified. Instead, find, identify, diagnose and prescribe the applicable treatment or form of data footprint reduction or other IT Infrastructure Resource Management (IRM)[1] technology, technique or best practices to cure the ailment.

Data footprint reduction (DFR) as a storage service is about storing more data in a denser footprint and not just for backup/restore. This includes storing more data managed per person where the additional data being retained adds value to an organization. Also included is keeping more data readily accessible, not necessarily instantly accessible, however within minutes instead of hours or days where access to more data adds value to the organization. Similar to storage tiering, the key theme with an effective DFR strategy is to avoid treating all data and applications the same. This means using several different tools, technologies and techniques to meet diverse application and data service level objectives (SLOs) requirements.

DFR is a collection of techniques, technology tools and best practices used to address data growth management challenges. Dedupe is a popular DFR technology with a market sweet spot for adoption being backup/restore or other highly repetitive data. However, the expanding scope of DFR is to cover primary and secondary active inactive (idle or dormant) data, some of which does not lend itself to dedupe, thus requiring different tools and techniques.

Common DFR technologies and techniques include:

- Archiving (structured database, semi-structured email, unstructured file or NAS data)
- Compression and compaction including real-time, streaming and posting processing
- Consolidation of storage and data centers, storage tiering and thin provisioning
- Data management including cleanup and deletion of unnecessary data
- Data deduplication or dedupe, also known as single instancing or normalization
- Spacing saving snapshots and replication optimization

The expanding focus of DFR pertains to active, On-line and primary data and storage as opposed to a single approach targeted at a specific area such as dedupe for backup/restore. Consequently the bigger and broader opportunity for DFR across organizations is to address different Performance, Availability, Capacity and Economic or energy efficiency (PACE) requirements using a variety of DFR techniques and technologies.

In other words, avoid developing tunnel vision on just one or a few techniques and missing important opportunities. This also means avoid trying to use just one tool to address all issues or challenges and, instead, align the applicable techniques and tools to the task at hand. If all you have is a hammer, then everything will look like a nail, but if you have a tool box along and a skill set on what tool to use for a given task, your efficiency and effectiveness are enabled for better productivity.

[1] IRM focuses on managing resources (servers, storage, and networks) including backup/restore, data protection, preservation, security, provisioning, reporting and capacity planning. Learn more in the book "The Green and Virtual Data Center" (CRC) as well as at www.storageio.com

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com

Data Footprint Reduction Innovation with Real-time compression

As previously mentioned, there are many different DFR approaches and technologies to address various storage capacity optimization needs. Likewise, there are different metrics to gauge the efficiency and effectiveness of the various approaches, some of which are time (performance) centric while others are space (capacity) focused.

Dedupe can be thought of as intelligent compression, although advocates of dedupe will probably disagree with that simple explanation. On the other hand, compression aims at some level of space or capacity DFR while maintaining a given data transfer or movement rate and is more closely aligned with performance or time sensitive needs. The use of dedupe or compression or some combination of the two is determined by the needs of the task at hand. As mentioned before, it's a matter of using the right tool for the job.

The different types of DFR technologies are also implemented in various ways in multiple locations. Real-time compression is focused on providing immediate DFR benefits instead of having to wait for deferred mode processing such as with some dedupe or post processing compression approaches. For active on-line and primary applications that are performance or time sensitive (where time means money) that require a DFR benefit, real-time compression is the tool from the IT storage services tool box to use.

In addition to performance or time sensitive applications that can benefit from DFR, real-time compression can also be applied to other applications or scenarios. For example, investment in current NAS storage solutions (hardware and software licensing) can be maximized by increasing the usage or usefulness of that solution by doing more work which reduces per unit cost while maintaining or perhaps enhancing performance. The result is that using real-time compression enables innovation to achieve objectives shown in Figure 2 (support growth; reduce cost without compromising QoS, SLOs or application performance).

Some of compression implementations are real-time, working on reducing or expanding data on the fly, with no impact to performance or applications. Other compression implementations may operate in-line for reads while writes are handled at a later time in the background on a post processing basis. Yet other compression techniques are time delayed for both reading and writing of data, resulting in no immediate DFR benefit for on-line storage or active applications.

Real-time compression should provide a DFR space saving benefit without negatively impacting reads or writes and be independent of the applications or types of data. True real-time compression should be transparent without performance penalties to applications while waiting for data to be compressed or decompressed. If you are looking to expand DFR focus from backup/restore and inactive data to a broader basis across your environment real-time compression should be considered. Innovation is supporting demand (more data) while reducing per unit costs without compromising QoS or missing SLOs.

A decision on DFR technologies should not be a binary either-or choice. Instead, take advantage of multiple available tools and use the right one for each situation. Additional DFR benefits can be achieved across various applications and storage tiers by combining tools and techniques.

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com

For example, compressing less frequently accessed data and then applying dedupe for Disk to Disk (D2D) backup/restore, business continuance (BC) and disaster recovery (DR) for additional savings or benefits.

Taking action, establishing a DFR strategy and what you can do today

- Develop a holistic approach to managing to your growing data footprint.
- Look beyond storage hardware costs; factor in software license and maintenance costs, power, cooling and IT staff management time.
- Leverage data compression as part of an overall data footprint reduction strategy.
- Optimize and leverage your investment in existing storage across all types of applications.
- For on-line, active, primary applications and storage, real-time compression provides immediate DFR benefits with transparency and without introducing performance bottlenecks.

An issue to consider is how much delay or resource consumption you can afford to use or lose to achieve a given level of data footprint reduction. For example, as you move from coarse (traditional compression) to granular, such as data de-duplication or single-instancing, more intelligence, processing power, or off-line post processing techniques are needed to look at larger patterns of data to eliminate duplication.

If you are evaluating other forms of data footprint reduction technologies for future use, including archiving with data discovery (indexing, discovery) or data de-duplication techniques, leverage appliance-based compression technology for immediate relief to maximize the effectiveness and capacity of existing storage resources for on-line, backup, and archive while complimenting other data footprint reduction capabilities.

The solution to one problem should not introduce issues or challenges elsewhere. Avoid introducing problems such as performance bottlenecks or increased workload on existing storage systems in the course of implementing DFR. As was seen in Figure 2, the objective is to reduce per unit costs (e.g. managed storage capacity) while supporting growth without compromising performance, QoS or missing SLOs. While it is possible to achieve DFR benefits with some solutions to drive costs down and support growth, care must be taken that it's not done at the expense of application performance impact. Consequently, look for DFR solutions that compliment your applications and storage environment providing investment protection without compromise.

Summary

Given that different applications and types of data along with associated storage mediums or tiers have various performance, availability, capacity, energy and economic characteristics, multiple data footprint impact reduction tools or techniques are needed.

What this means is that the focus of data footprint reduction (DFR) is expanding beyond that of just dedupe for backup or other early deployment scenarios. For some applications reduction ratios are an important focus so the need is for tools or modes of operation that achieve those desired results. For other applications the focus is on performance with some data reduction benefit, tools are optimized for performance first and reduction secondary.

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com

General action items include:

- Develop a DFR strategy for On/off-line data that is complementary to each one another
- Measure and compare storage based on idle and active workload conditions
- Storage efficiency metrics include IOPS or bandwidth per watt for active data
- Small percentage reductions on a large scale have big benefits
- Align the applicable DFR technique and technology to the task at hand
- Dedupe is good when there is high-degree of redundant data (VMs and Backup)
- Combine various DFR techniques and technologies with different best practices
- Balance data reduction ratios (space) with the need for data transfer rates (time)
- Innovation is aligning cost saving, supporting growth without compromising QoS or SLO
- Look at real-time compression for a balance of performance and DFR for active data
- Maximize investment in your existing storage by using real-time compression
- On-line active and primary applications need real-time compression storage services

In short, deploy a comprehensive data footprint reduction strategy combining various techniques and technologies to address point solution needs as well as your overall environment, including on-line, near-line for backup, and off-line for archive data. If you need immediate DFR benefits for both reading and writing data while maintaining application transparency without introduction of performance bottlenecks, then real-time compression is your best choice. The bottom line is to use the most applicable technologies or combination of technologies along with best practices for the task and activity at hand. Learn more at www.storageio.com or www.storageioblog.com as well as at IBM Real-time compression landing page which is located at www.realtimecompression.com.

About the author

Greg Schulz is founder of Server and StorageIO, an IT industry advisory consultancy firm, and author of the books *The Green and Virtual Data Center* (CRC) and *Resilient Storage Network* (Elsevier). Learn more at www.storageio.com, www.storageioblog.com or on twitter @storageio.