TANEJA GROUP
TECHNOLOGY ANALYSTS

T E C H N O L O G Y B R I E F

Matching Data Deduplication Approaches to Workload Characteristics: The Quantum DXi7500

December 2008

Storage capacity optimization (SCO) offers undeniable economic benefits to enterprises of all sizes by significantly reducing the amount of raw storage capacity required to store any given amount of information. The use of SCO technologies drives savings not only in hardware and software purchases, but also in cooling, floor space, administration, and network capacity (bandwidth) expense. The most commonly deployed SCO technology is data deduplication, and there has been considerable controversy among vendors about the "best" approach to deploying deduplication technology. In our opinion, there is no single "best" approach to data deduplication; it is important to determine the problem that needs to be solved, and then select the deployment mode that best meets those requirements.

The primary deployment environment today for data deduplication is secondary storage. Secondary storage applications such as backup and archive have characteristics, such as medium to high latency of access requirements and significant redundancy, which make them an excellent target for deduplication. In this Technology Brief, we'll discuss different approaches to data deduplication in the market today along with the pros and cons of each, and then match some of the most common backup problems to a recommended deduplication approach. We'll close with a spotlight on Quantum, a vendor whose technology is unique in offering a single solution with multiple concurrent data deduplication modes that can be matched to different workloads. This "policy-based" approach, as Quantum calls it, is a key feature of the company's DXi7500 enterprise disk backup and replication system. The flexibility it provides makes the DXi7500 a compelling choice for users that are trying to address a number of different backup problems across mixed workloads – a description that fits most enterprises – with a single solution.

Defining Storage Capacity Optimization (SCO)

SCO technologies basically use one or more algorithms to reduce the amount of raw storage capacity required to store any given amount of information. There are technologies available today targeted for use with both primary and secondary storage applications. Primary and secondary data have very different characteristics, particularly in terms of access latency requirements and redundancy. A complete discussion of these differences is beyond the scope of this paper, but it is important to note that the capacity optimization ratios achievable against secondary storage tend to be significantly higher than those achievable against primary storage. SCO technologies include enhanced compression, file-level

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

single-instancing, and block-level deduplication, among others. Because they are most commonly used today in backup environments, we will provide a distinct data protection focus to our discussion.

Defining Data Deduplication

At their most basic level, data deduplication algorithms break data down into recognizable "chunks" and then look for redundant chunks. When a redundant chunk is found, that chunk is replaced by a pointer to a reference chunk, and the full chunk is deleted. In secondary storage environments like backup where backed up data may only change 3-5% or less per day, there is a significant amount of redundancy that can be identified and removed. In contrast to single-instancing technologies which remove redundancies at the file level, deduplication algorithms operate at the block level, a lower level of granularity which generally produces higher capacity optimization ratios than single-instancing. Chunks can also be of fixed or variable lengths, with variable length chunking having an edge in achieving higher capacity optimization ratios. Looking across vendor offerings in the market today, it is not unreasonable to reduce the disk needed to store backup data by 95% or more over time.

Deduplication Approaches

Deduplication can be performed either at the source or at the target. With source-based deduplication, data is deduplicated, using backup client (application server) resources, before it is sent across a network. With target-based deduplication, the data is not deduplicated until it hits the storage target. Source-based approaches can reduce the amount of network bandwidth required to complete initial backups, but can put unwanted additional load on an application server and tend to only be able to handle smaller data volumes. When application server overhead is an issue or large data volumes are present, as in most centralized data protection environments, target-based deduplication can offload the application server(s) and generally offer much higher deduplication processing speeds.

Source-based deduplication technologies all operate in-line (in-band) but target-based deduplication can be done either in-line (during the initial backup ingest) or through a post-processing approach which is asynchronous and subsequent to the initial backup ingest. In-line approaches capacity optimize the data before it ever hits the storage target, resulting in immediate storage capacity savings. In-line approaches can, however, impact backup performance. Basic laws of physics dictate that an in-line approach must introduce some additional latency, although depending on the performance requirements of a particular environment, the performance impact may not be noticeable. If the deduplication process cannot be performed at wire speeds though, there is the possibility that it can result in longer initial backup windows.

Post-processing deduplication is an out-of-band approach where data is not de-duplicated until after the backup has completed. The advantage to this approach is that the backup window will never be negatively affected as a result of

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

deduplication. There are two disadvantages. First, because it does not reduce backup "landing pad" raw storage requirements at all, you will need more overall storage capacity relative to in-line approaches. And second, because backup ingest and data deduplication are two separate processes performed consecutively, they can take significantly more time to get the data into its "capacity optimized" state and ready for further processing, such as replication to a remote site. Think of post-processing as a "trash compactor" for your data: once the backup is completed, the data deduplication process begins, and reduces the amount of storage required to store that particular backup going forward.

The Positive Implications of Data Deduplication

Reducing the amount of data required to store any given information has a number of positive implications. First, and most evident, it can reduce the amount of raw storage capacity required, and hence reduce the spend not only on storage hardware but also on the power, cooling, and floorspace required. Second, the fact that you're now dealing with less data means that you do not need as much bandwidth to move it to another location. If backups, upon completion, are replicated to a remote site for disaster recovery (DR) purposes, data deduplication can significantly reduce the amount of data that has to be sent, allowing these large sequential data transfers to complete much more rapidly or with significantly less bandwidth consumption. In either case, using replication instead of physically shipping tapes to remote sites for DR purposes can have a very positive impact on DR recovery point objective (RPO).

Define Your Problem Correctly

In selecting the right deduplication method, it is important to first determine which backup problem is most important to solve. Do you want the shortest backup window? Do you want to minimize raw storage capacity requirements? Or do you want the fastest time to replication (in cases where a backup is being replicated between sites for DR purposes) for better DR RPOs? For jobs that require the shortest backup window, a post-processing approach to deduplication can be most appropriate. When the desire to minimize raw storage capacity requirements is paramount, in-line deduplication can be a better approach. For data that doesn't de-duplicate well or won't be retained, then it may be appropriate not to use data deduplication at all.

What is interesting to note is that most companies have a mix of jobs that demand different requirements. For smaller companies or very limited environments, backup may be viewed monolithically, but most companies have different backup jobs with different requirements—some need the shortest backup times, while others would benefit more from higher rates of capacity savings. For certain data sets that offer very low levels of redundancy, or have very high read performance or stringent recovery requirements, it may be better not to apply any deduplication at all.

Selecting the right data deduplication approach depends, then, upon what problem

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

you want to solve. Vendors that offer only in-line deduplication may offer the optimal solution as long as the only metric you want to optimize for is minimizing raw storage capacity requirements, while vendors with a post-processing approach can guarantee zero impact on backup windows at the expense of raw storage capacity. A vendor, however, that offers a data deduplication solution that supports multiple approaches, allowing an administrator to apply the approach by job type, may offer the best solution for environments with mixed workloads and especially across distributed organizations that may have varied backup servers, applications and backup methodologies.

Quantum's Policy-Based Deduplication Approach

Quantum entered the SCO market in early 2007 with the introduction of backup appliances (the DXi3500 and DXi5500) which used a unique new approach to data deduplication called "adaptive mode". In adaptive mode, deduplication occurs during backup ingestion but it manages the data flow by writing the data to a disk buffer in small segments. As soon as a segment has landed, it is then picked up and fed into the deduplication process. As long as deduplication speed keeps up with the speed at which the backup is ingesting, adaptive deduplication appears to operate in an in-line manner. If backup data is being ingested faster than it can be deduplicated, for example when there is an input burst, the disk buffer will start to fill up; in the meantime, deduplication continues to operate on segments that have already been written into the buffer. Disk buffers can be configured to be of any size. With adaptive mode, backup ingestion and deduplication are occurring simultaneously, so even though they are two separate processes, the backup data is converted into capacity optimized form much faster than if it was converted using a pure post-processing method which did not start until the entire backup had completed. Buffering the backup data stream has the advantage of enabling the system to adapt to input bursts, ensuring that backup performance will not be affected even when deduplication processing performance is exceeded.

In designing its next generation Dxi7500 solution, Quantum recognized that enterprise customers would want to optimize for different data sets and backup jobs. As a result, they developed "policy-based deduplication," which gives users three different deduplication options (or modes) in a single backup appliance and the flexibility to deploy the most appropriate one for each job to best meet their disk capacity or backup/recovery window requirements. All Quantum deduplication modes operate at the block-level and use variable length windows for higher capacity optimization ratios.

Adaptive Mode

As explained above, the adaptive mode effectively operates like a conventional in-line approach unless the backup ingest rate outruns deduplication processing speed. At this point, the deduplication process adjusts so that backup performance is not affected. The benefit of Quantum's adaptive approach is that it can maintain very high backup performance (higher than that available with conventional in-line approaches) while

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

introducing no more than a minimal lag time between the completion of the backup ingest and the data deduplication processes. For larger backups, the adaptive approach may actually complete both the backup ingestion and the deduplication process faster than pure in-line approaches (which tend to support much lower backup ingest rates). Adaptive deduplication is recommended for workloads that have the following characteristics or requirements:

- A medium to short (but not the shortest possible) backup window
- A balanced trade-off between high performance (short initial backup window) and minimized raw storage capacity
- A remote site replication configuration where remote site RPO is very important
- A setting with backup jobs exhibiting bursty I/O that varies by job

Remote Office/Branch Office (ROBO) environments that require a local backup appliance tend to be a good fit for adaptive deduplication, as do environments with a combination of both physical and virtual server deployments. Mailbox backups where compound documents (attachments) are involved also tend to be a good match for adaptive deduplication.

Post-Processing Mode

Quantum's post-processing mode lets the backup data land on disk, deferring deduplication activity until all of the data has been collected or until a scheduled time period. Because all the system processing is devoted to ingest, the deferred mode offers the shortest backup window while still significantly decreasing the raw disk capacity needed to store backups over time. This mode of deduplication is recommended for environments that have the following characteristics or requirements:

- Maintaining the shortest possible backup window is paramount
- Relatively higher percentages of unique data than other workloads
- Some additional disk space can be allocated for temporary use

Environments with larger amounts of unique (as opposed to redundant) data tend to be a better fit for post-processing deduplication, as do large on-line transaction processing (OLTP) database environments. More broadly, any environments which are trying to significantly cut down on the costs associated with backup storage while at the same time maintaining the shortest possible backup windows tend also to be good candidates.

Native NAS/VTL Mode

The native NAS/VTL mode delivers the fastest possible performance by offering conventional disk backup without deduplication. This mode can also be thought of as a "pass through" mode.

Although it may seem strange that this is listed as a deduplication "option", there are certain workloads for which deduplication is just not appropriate. In cases where workloads have very stringent read requirements that cannot be met when data must be converted back to its original form, the use of deduplication is not recommended. This may become even more of an issue in low

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

latency environments where very large restores are performed. There may also be certain types of data that do not deduplicate very well, leading to very low (<3:1) data reduction ratios. In these cases, it may not be worth performing the deduplication process against this data and more appropriate to just use standard compression approaches.

In mixed environments with backup jobs and data types across all three of these areas, it can be important to be able to selectively deploy different data deduplication approaches against different data sets. Solutions which support only one approach can be limiting in these mixed workload environments whereas solutions which support multiple approaches simultaneously against different data sets can offer a more tailored solution.

Quantum's DXi7500

Quantum has incorporated its policy-based deduplication technology into the DXi7500, its enterprise disk backup and replication appliance introduced earlier this year.

Like Quantum's DXi3500 and DXi5500 platforms, the DXi7500 can be configured to present either a NAS or a FC (VTL) interface. As an appliance, it front ends a FC storage area network (SAN) where the deduplication repository resides. Significant redundancy has been built-in to the DXi7500, including dual RAID controllers with active/active failover, redundant power, cooling and fans, and hot replaceable components. Single node configurations can process up to 4TB/hour of data. With their DXi3500 and DXi5500 offerings, Quantum offered viable adaptive deduplication solutions for the mid market. With the introduction of the DXi7500, Quantum can now scale to configurations that support up to 180TB of usable capacity (the amount of storage after RAID or mirroring overhead has been taken into account), expanding soon to 220TB, and almost 4PB of deduplicated capacity (depending on achieved capacity optimization ratios).

Integrated tape support on the DXi7500 allows customers to consolidate and centralize physical tape creation from disk-based sources, minimizing the need for sophisticated administrative resources in remote locations. Optional replication capabilities allow DXi7500 systems to replicate backup jobs asynchronously across IP-based networks, enabling the easy creation and maintenance of replicated copies for DR purposes. Support for any backup software, as well as common management across disk, (virtual) tape, and replication – all of which leverage a global deduplication repository for higher capacity optimization ratios – make the DXi7500 an anchor platform for enterprise backup. In addition, with the ability to replicate data to and from a DXi3500 or DXi5500 in a ROBO or intermediate data center, the DXi7500 provides an edge-to-core backup and recovery solution in a single, scalable architecture.

All Quantum's deduplication products support adaptive deduplication, but only the DXi7500 adds support for job-specific, policy-based deduplication. Mixed workloads running on different servers can

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

use it as a backup target, with the optimal approach to deduplication – adaptive, post-processing, or native NAS/VTL (no deduplication) – deployed to meet specific backup job parameters and requirements simultaneously.

Taneja Group Opinion

Quantum's choice to allow customers more than one deduplication approach in a single product is unique in the industry. This provides the flexibility to deal with the mixed backup workloads that are so common in most enterprises, configuring the best deduplication policy for each as they all flow through the same appliance.

Quantum's simultaneous multi-mode support is not their only differentiating technology though. The ability of Quantum's adaptive deduplication to handle bursty and other varying and unpredictable I/O requirements while minimizing the time to capacity optimize the backup data is an excellent match for the varying workloads found in most enterprise environments. And Quantum's DXi product portfolio offers customers a complete line of cost-effective edge-to-core solutions that comprehensively cover data protection and DR requirements while minimizing network bandwidth consumption.

Quantum's innovative approach in combining multiple deduplication methods clearly offers more flexibility in mixed workload environments than other single solutions. In our opinion, most enterprises do have mixed workloads, and Quantum's approach is meaningfully better in all but the most monolithic environments. We see other major vendors agreeing in principle with Quantum that different types of environments require different deduplication implementations, and targeting their products towards workloads for which they are best suited. Only Quantum, EMC, and soon Dell (through their licensing of Quantum's deduplication and replication software) offer this measure of flexibility in a single product. That flexibility translates to a solution that can be tuned to simultaneously meet the needs of varying workloads better than any "pure play" solution (in-line or post-processing). If you are looking to leverage data deduplication technologies against mixed workloads to lower your overall storage costs, Quantum offers a sustainably differentiated solution in a single product.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com