

THE ROI OF REAL-TIME COMPRESSION



Get Instant ROI, while you plan the rest of your Primary Storage Optimization Strategy

2009 is going to be a year of cost containment and nowhere is the opportunity to contain costs more apparent than in primary storage. There are many techniques available now to reduce the capacity demands of primary storage like archiving, thin provisioning and deduplication but probably none delivers the rapid and immediate ROI as does real-time compression technologies like those from Storwize Inc.

These technologies work with your existing storage infrastructure to optimize primary disk storage. Storwize for example will work with any NFS or CIFS based NAS or file server. This makes integration easy and fast. The technology lives on a dedicated appliance between the users and NAS. All data that passes through the appliance is compressed in real-time before it is stored.

While results may vary it is typically safe to expect at least 50% reduction in primary storage utilization by using real-time data compression. Some environments, depending on the compressibility of the data set, may have as much as 80 or even 90% reduction. For data centers considering a purchase of additional NAS based storage in 2009 this can certainly reduce the size of that expenditure and for many customers eliminate it.

Real-time compression delivers a unique and compelling ROI. A key difference between real-time data compression

and other technologies is that it is effective on ALL the data stored on the NAS environment. This can be highly active data like Oracle or VMware, or it can be data that is aging and may be in fact ready for archiving. Reducing data at the point of origin is unique to real-time compression and delivers ROI not only in the reduction of the storage required by active data but also all the capacity requirements of the copies of active data are reduced.

For example one of the key limiters in the number of snapshots that can be maintained for a NAS environment is the size of the snapshot reserve area. This area holds all the unique writes between the snapshot data and the live copy of data. Real-time compression improves the efficiency of the snapshot reserve area allowing the retention of twice or more as many snapshot copies of data.

In another example, backups and WAN replication become instantly more efficient when real-time data compression is used. Because there is less data to backup or transmit less network bandwidth is required and less storage (disk backup, tape or replicated disk) is required to store that secondary copy of data. Organizations can also deploy real-time compression appliances, or optional client software, in remote sites to insure quick recovery of the compressed data.

Another ROI component unique to real-time compression is as a compliment to any eventual archive strategy. Because data is reduced at primary, when moving that data to the secondary disk archive, the amount of data that has to be moved and stored by the archive device is substantially smaller. This reduces purchase costs and slows the growth of the disk archive area.

Real-Time compression is also complimentary to other data reduction techniques in addition to optimizing the move to an archive. For example an increasing number of NAS vendors are offering either deduplication or single instance storage. Compression used in conjunction with deduplication has proven enhance the overall space savings when compared to deduplication alone.

The final and possibly most important area of ROI benefit is that real time data compression tools are simple to implement and take advantage of from day one. Unlike an archive project that requires a decision on what software to use, what data to archive, how long to retain it and how and where to store it, real-time compression simply inserts between your users and the NAS appliance and starts working. No decisions need to be made, no

business process to change, it just starts working.

Real-time compression vendors typically get asked about performance but they have a good answer.

Independent test results have shown that in most cases users will see an actual increase in performance and it's quite simple to understand. With real-time compression the amount of data that has to pass through the storage infrastructure is effectively cut by 50% or more. This makes your storage bandwidth twice as fast, your network bandwidth able to transfer twice as much data and your storage controller cache memory able to hold twice as much data. The results are impressive even in such IO intensive applications like NFS mounted Oracle databases as Storage Switzerland covered in our article on Data Reducing Oracle, or in highly random IO environments like VMware as we covered in our article for SearchServerVirtualization.com.

Another typical question is about data integrity. These systems now offer end-to-end data verification, fault detection and correction throughout the entire data path, including verification of the ability to read the compressed storage. This insures the ability to read the data. In addition to highly redundant and available

systems, leading solutions support high-availability configurations and include client software for reading data in remote locations where an appliance is not available.

The result is a solution that delivers a near instant return on investment. The product can just start compressing data as its accessed. In addition, most implementations take advantage of the ability to scan the existing file system, compressing everything in advance for optimal space savings.

About Storage Switzerland

Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. Storage Switzerland provides strategic consulting and analysis to storage users, suppliers and integrators.