Simple, Reliable, Fast De-duplication, Data Protection Platform

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) White Paper
Prepared for IBM

October 2008

EMA

Table of Contents

Executive Summary 1

Introduction 1

The Case for VTLs 1

Data De-duplication 2

Why a Clustered Gateway? 3

The IBM System Storage TS7650G ProtecTIER De-duplication Gateway 4

The EMA Perspective 5

About Diligent, an IBM company 8

- About Diligent Technologies: 8
- About IBM: 8

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

Simple, Reliable, Fast De-duplication, Data Protection Platform

Executive Summary

While Virtual Tape Libraries (VTLs) are now a fundamental element of the data center infrastructure, the growth of data under management is outpacing most firms' ability to add disk capacity to their VTLs. This phenomenon has led to a strong demand for data de-duplication. Data de-duplication is a technology that finds and eliminates redundant data within a disk repository. The effect is a dramatic increase in the usable capacity of a given disk pool. Combining this powerful technology with VTLs allows firms to store and retain far more data than they otherwise could, while saving significantly on disk storage costs.

However, to meet requirements of the enterprise data center, a VTL must do more than simply reduce storage requirements. To be viable, it must maintain high sustained throughput performance, and must meet the availability demands of the enterprise.

The newly released TS7650G de-duplication gateway from IBM is an example of a de-duplicating VTL that meets the demands of the enterprise data center. It is the first solution that supports a cluster of two nodes while sustaining high performance, inline de-duplication.

Introduction

While concerns over regulatory compliance, e-discovery, and the continuing data explosion add to the pressures that IT managers cope with, the IT manager's job remains essentially the same as it has been for the last 30 years in one important respect: at the end of the day, the IT team's measurement of success has everything to do with how well they answer the question "is the company's data protected to the degree that it must be, and are we assured of business continuity no matter what the conditions?"

Protecting business-critical data remains job one, whether a data center measures its business-critical information in gigabytes or petabytes. But keeping pace with new demands while at the same time coping with massive quantities of new data poses a significant challenge. New backup and recovery technologies barely keep pace with the need to protect and recover data quickly, and long-term archives must now accommodate the requirements of regulatory compliance, the edicts of corporate governance, legal discovery demands, and the needs that many business units now have for quick access to historical data to support complex business analyses.

This document looks at how a virtual tape library (VTL) utilizing a clustered gateway and data de-duplication can assist data centers that face such challenges, and at how the introduction of such devices provides data center managers with an easy route to faster backups and recoveries, more efficient use of existing hardware, and simplified management of even petabyte-sized data repositories.

The Case for VTLs

Virtual tape libraries are used in the same way that physical tape devices would be used. The VTL is inserted as a staging area between the first-line disk devices and the tape libraries, and from then on it appears to the backup and recovery software to be a physical tape library. This changes the physical paradigm for backups from disk-to-tape to disk-to-disk-to-tape (D2D2T), a shift that has some interesting consequences.

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

All VTLs perform the same basic functions, serving as a disk-based target for the backup process, and offering rapid recoveries from local disk-based storage. That being said, however, all VTLs are certainly not equal. Users of first-generation VTLs, which use a disk-to-disk (D2D) backup model, found that as their data protection needs increased they ran out of disk capacity very quickly. The ability of second-generation VTLs to migrate older data to tape (the D2D2T model) makes for better economic sense, and allows managers to set their own policies regarding trade-offs between disk and tape. Additionally, many VTLs provide value through additional features such as encryption, compression, and data de-duplication. All are potentially useful, although de-duplication offers the clearest evidence for rapid payback on the VTL investment.

Users of VTLs should expect the following:

- Maintenance windows will be significantly reduced, with the result that production systems stay online longer.
- Backups are CPU-intensive, as backups to tape take place in the VTL; however, performance degradation of production systems will be sharply reduced.
- Because recoveries will be far faster, downtime during a recovery will be sharply reduced.
- Because incremental backups no longer require tape media, IT rooms immediately lower their expenses associated with tape media and tape handling.
- Perhaps most significantly, because the system appears unchanged to the backup and recovery software IT departments can continue to use the same backup scripts and policy-based management techniques that were in use prior to VTL's arrival.
- As the uptime for business processes increases, service-level objectives (RPOs and RTOs) are more easily achievable, stakeholder satisfaction should increase.

Data De-duplication

"Data de-duplication" refers to a set of technologies whose goal is to reduce the total amount of data that must be managed while maintaining the integrity, availability, and accessibility of that data. By reducing the amount of data under management, data centers reduce the amount of hardware required – thus decreasing operational expenses for power and cooling – reduce tape media costs, and make it easier for storage admins to cope with drastic increases in raw data and increased complexity in the rest of their storage infrastructures.

By reducing the amount of data under management, data centers reduce the amount of hardware required.

Vendors implement de-duplication in several different ways. The main issues to consider are:

- **Performance**. No matter how effective a technology is at reducing capacity requirements, it will not be viable unless data can be processed within the given time window (thereby meeting various service level agreements for backup completion and off-site data protection). Performance is always a major criterion when analyzing de-duplication solutions.

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

- **Granularity**. De-duplication can be done at the file-, block-, or byte-level. The more granular the analysis, the more efficient the process will be.
- **Factoring** (identifying redundant data). De-duplication relies on proprietary algorithms, based either on hashing (which first creates a "fingerprint" for each block of data, and then compares the fingerprints), or on comparing actual data. Factoring is the most important determinant of de-duplication speed.
- **Architecture**. Deduping architectures fall into one of three categories.
  - *Parallel processing* writes the raw data to disk, de-duplicates it, and rewrites the factored data at the same time. Running these three operations in parallel requires both extra processing power and additional disk space.
  - *Post-processing* first writes the raw data to disk, de-duplicates the data, rewrites the de-duplicated data, and finally erases the raw data. This multi-step procedure requires less processing power because it defers the factoring, but it also requires extra disk space and takes longer, threatening backup windows.
  - *In-line processing* de-duplicates the data before it is written to disk. This requires the most processing power, but makes the most efficient use of disk space and backup windows.
- **Where the process takes place**. De-duplication can be run on either hosts or virtual tape libraries. When run on a host, this will require reconfiguring the host environment, frequently resulting in a need for forklift upgrades and re-scripting and testing of existing backup procedures. When processing takes place on a VTL, however, de-duplication is transparent to existing processes; new functionality does not interfere with existing best practices.

The two major metrics used to describe de-duplication are *throughput* and *factoring ratio*. Throughput is a measure of the time it takes to write the factored data to the VTL's storage. The factoring (or de-duplication) ratio is derived by dividing the total data before reduction into the total data after reduction. Readers are advised to make every effort to understand how vendors use these figures as a number of claims may be suspect.

Why a Clustered Gateway?

A storage gateway connects backup arrays on one side of the gateway to backup servers on the other (see Figure 1). A key part of the value proposition of such devices is that, because a gateway comes with no storage itself, storage managers can assemble combinations of devices on either or both sides of the gateway from a variety of vendors and device types. The nonproprietary nature of gateway-based storage means that managers can *preserve relationships with their current vendors*, and *can continue to extract value from existing storage hardware*. When it's time to scale out (or scale up) with new arrays, managers are typically free to stay with their original set of vendors or to explore relationships with others. Whatever the choice, gateways make it easy to add additional storage when required.

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

A clustered gateway can help provide high availability.

Enterprises value clustered storage gateways for the same reasons they value clustered application servers. Clustering can scale *performance*, which at many sites will be the crucial differentiator when it is time to purchase. A clustered gateway can help provide *high availability* – should one node in the cluster fail, failover to the second node can ensure that operations continue until the failed node is brought back online. Thus, as with clustered servers, clustered storage gateways can support an IT strategy that requires eliminating any single point of failure within a corporate data center.

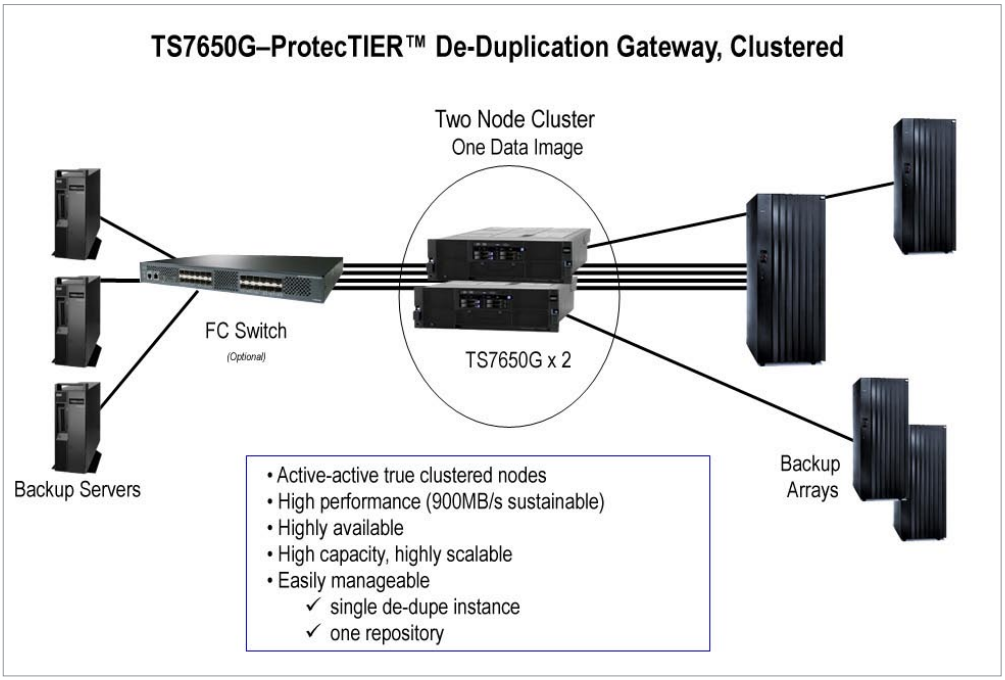

Figure 1 The clustered storage gateway.

The IBM System Storage TS7650G ProtecTIER De-duplication Gateway

Based on Diligent's industry-proven ProtecTIER, the IBM TS7650G is a VTL gateway that resides on a standard 4U Linux server. It is deployed between the back-end storage arrays and the backup media-servers, and can accommodate a petabyte of raw data. The gateway comes standard with four Quad Core Xeon processors, 32 gigabytes of RAM, dual gigabit Ethernet ports with TCP/IP offload engines, and four dual port PCI express 4 GB Fibre Channel cards. Each node is capable of supporting up to 1 petabyte of physical capacity, and comes with IBM's patented HyperFactor data de-duplication engine. Disk support is available for a wide variety of array choices including IBM DS4000, DS8000, and XIV, as well as for non-IBM disk devices.

The device is available singly and in two-node clusters. Clustered configurations are active-active; this provides high performance/throughput and rapid semi-automated failover and failback capability, minimizing cluster downtime. Should a node fail, clustering allows

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

full repository sharing between the nodes – each node is capable of reading and writing to the data repository and each has access to all virtual devices. The TS7650G gateway enables high sustained *inline* throughput, scaling at a rate of over 450 MB/s per node. A two-node cluster therefore ingests and de-duplicates data at rates exceeding 900 MB/s in real customer environments. No other vendor offering inline de-duplication can approach these speeds.

The HyperFactor De-duplication Engine

IBM's HyperFactor engine uses proprietary algorithms developed by Diligent Technologies for its ProtecTIER VTL to perform high-speed pattern matching, identifying, and eliminating instances of duplicate data. This can reduce the total amount of stored data (and thus, the need to buy additional storage) by up to 25:1, depending on the environment. HyperFactor stands apart from other de-duplication technologies in several important ways. First, because the entire HyperFactor index is memory resident, the process delivers sustainable enterprise-scale throughput of over 450 MB/s per node. Hash-based technologies must make frequent disk reads, and so are either limited to inferior in-band transfer speeds or rely on out-of-band post-processing which is both costly and cumbersome. Secondly, whereas other de-duplication techniques impose severe capacity restraints, HyperFactor can accommodate data repositories of up to one physical petabyte. Finally, because HyperFactor filters out redundant data only after performing byte-level comparisons of the new and existing data, HyperFactor can support enterprise class data integrity. Other methods compare hashing codes rather than data, and may incur some level of risk.

The EMA Perspective

De-duplication. Data de-duplication can be a big cost saver when it comes to purchasing hardware and tape media, and a reduction in the overall mass of data being managed can reduce the day-to-day expenses associated with data protection. But de-duplication techniques vary significantly between vendors, and some offer significantly more flexibility and potential value than do others. Each of the three de-duplication technologies described in this paper – in-line, parallel, and post-processing – has its own advocates, and many vendors claim aggressively high numbers when describing their product's de-duplication ratios and throughput rates. Inline de-duplication has the advantage of completing the de-duplication task at the same time as the data ingest. VTLs that can sustain very high inline de-duplication rates will be attractive to a large set of enterprise customers. The ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) team recommends that anyone considering an investment in de-duplication technology understands thoroughly what a vendor's performance numbers mean, and how those numbers were derived. (See the accompanying sidebar.)

Gateways and Clustering. EMA feels that a properly designed storage gateway delivers value in two ways. First, because these devices can work with many standards-compliant storage systems, existing investment in storage hardware can be protected. As a result, when existing arrays and JBODs are cabled behind the gateway, they continue to provide value for as long as they are kept operational. Secondly, when new storage devices are required, devices from a variety of vendors can be used, allowing vendor selection to become as

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

much a business decision as it is a technology choice. IT managers will be free to pick and choose whichever vendor they feel most comfortable with, can play one vendor off against another, and can squeeze more value from their storage capital expenditures.

Clustered gateways are a high availability improvement over standard gateways, providing both high performance and scale-out capability while at the same time enhancing business continuity. The reader is advised, however, that all clustered gateways are not equal. The rapid failover capability that active-active configurations have provides a level of protection that active-passive systems typically cannot offer. In the case of the system described in this paper, failover and failback (and scale-out functions) are accomplished with little or no operator intervention. The level of protection that automated failover/failback delivers means continuous data availability and insured uptime for critical business processes.

A Caveat: The De-duplication Benchmarking Game

Obviously, the greater the efficiency of the de-duplication process, the greater the value of a de-duplication product will be. But what exactly is "efficiency" when it comes to de-duplication?

De-duplication essentially consists of two procedures. In the first, data is ingested into the process (this means bringing the data on board the VTL or other device where the de-duplication takes place). During the second phase, the ingested data is compared against the existing data corpus, duplicated data is identified, and the non-duplicated data is written to disk.

On the surface, it might appear that the de-duplication process is easy to measure, and in fact it is. But the numbers can be made to lie. Vendors benchmark their processes according to their products' strengths, and quote throughput numbers and de-duplication ratios separately. This can be quite legitimate, but the presentation can be deceptive. While some vendors measure throughput as the total time taken for both the ingestion and de-duplication process, others do not. Some vendors, most frequently those who use post-process de-duplication, consider throughput to be the speed at which data is taken into the process and written to the VTL. They do not, however, include the actual time required to de-duplicate the data, which in any postprocessing technique only begins once all the data has been moved to the VTL's disks.

Buyers should keep in mind that when it comes to data de-duplication, a significant metric is "time to the desired outcome," which is to say, the total time it takes to deliver a fully factored set of data to disk.

EMA feels that, particularly where throughput numbers are concerned, the only legitimate comparisons are those that compare one finished product with another. After all, the purpose of an investment in de-duplication technology is... to de-duplicate data. It makes little sense then to benchmark only part of that process. Buyers should keep in mind that when it comes to data de-duplication, a significant metric is "time to the desired outcome," which is to say, the total time it takes to deliver a fully factored set of data to disk.

Value from virtual tape libraries. EMA has always been a big fan of VTLs because they can provide a significant and easily calculated return on investment. VTLs are proven technology

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

whose best-known benefits are that they speed up both the backup and recovery process, they free up production devices from the heavy performance penalties imposed by supporting backups to physical tape, they reduce the number of tapes that must be handled, and they are relatively easy to bring online without adding additional IT overhead. VTLs offer another benefit however, one that dramatically increases in value as sites get larger.

Enterprise data centers invariably have a significant investment in their backup and recovery processes, going to great lengths to automate these processes in order to save on the daily operational expenses associated with backup and archiving. Scripted or policy-based, these processes are the sum of IT knowledge at each site, and represent a site's understanding of what its own best practices are. The largest data centers may have developed processes over decades, and their investment in designing, testing, and maintaining the scripts that drive those processes cannot be overstated. Because a virtual tape library looks to all processes just like the physical tape library it represents, any investment in scripts and policies is protected. When a VTL is inserted into the data center, a key change in the processes is that backups and retrievals can happen more quickly because data is now written to disk rather than physical tape. Physical tape can still be the initial backup target for data that will have infrequent access. Physical tape remains an important data retention tier in the storage hierarchy, addressing longer term data retention needs where TCO and energy conservation have priority over other considerations.

The IBM TS7650G, ProtecTIER De-duplication Gateway. EMA has closely tracked both the VTL and de-duplication markets since their inception. We recommend VTLs to our IT clients as an easy way to preserve backup windows and enable very fast data recovery, and recommend de-duplication as the most effective data reduction technique available. We first began looking at the HyperFactor de-duplication technology when it was released by Diligent Technologies several years ago as an integral part of their ProtecTIER and VTL. We were impressed with its speed and efficiency then, and are more impressed with it now. Sustained throughput rates currently exceed 450 MB/second (per node) in real customer environments; adding a second node (clustering) doubles this number, scaling both throughput and capacity, and adds high-availability. If IBM (which purchased Diligent Technologies in early 2008) can increase the clustering capabilities beyond two nodes, we will be happier still. (ProtecTIER architecture supports future expansion)

the TS7650G provides excellent de-duplication, and does it with stellar throughput rates.

That being said, the TS7650G provides excellent de-duplication, and does it with stellar throughput rates. This is a direct result of the HyperFactor technology, which is a clear differentiator for this product. Hashing technologies requires frequent disk reads during the factoring process, a performance bottleneck. HyperFactor, however, does not use hashing; all required metadata is maintained in the TS7650G's 4 GB of RAM no matter what the size of the data repository. Because the factoring index is accessed in microseconds (rather than the milliseconds required for a disk read), de-duplication throughput is much faster than competitive technologies can offer.

Storage administrators are also likely to appreciate the way the TS7650G makes it easy to manage large amounts of backup data. All de-duplicated data behind the gateway (as much

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

as a petabyte) is seen as a single target for all policies and media servers; the scalability of the system means that recent backups can be kept online according to the dictates of local IT policy, and that recoveries – majority of which will come from the most recent set of backup data – can be accessible at very high speeds. Furthermore, HyperFactor's data reduction makes it practical to transfer backed up or replicated data to remote sites electronically. Planning for disaster recovery becomes very practical.

Making a business case. Managers looking to invest in de-duplication technology should have little trouble making a business case to justify an investment in the TS7650G. The TS7650G with HyperFactor:

- Defers the need for short-term hardware investment because HyperFactor de-duplication can free up space on existing arrays.
- Reduces the long-term need for additional disk storage, as only changed data blocks are captured when new data is written to the VTL.
- Positively affects budgets for power and cooling by deferring the purchase of added disk hardware.
- Allows vast quantities of data to be kept on hand and readily accessible. It is thus easily referenceable should corporate officers be required to respond to internal governance or regulatory compliance issues.
- Can support up to 1 petabyte of physical storage capacity per virtualization engine, which represents up to 25 petabytes of de-duplicated capacity.
- Is capable of sustainable inline de-duplication throughput of over 900MB/s per two cluster node
- Optimizes processing requirements because the product uses inline de-duplication. That is, the redundant data is reduced as it is ingested. This eliminates the need for a dedicated "post-process" window that is required by many other de-duplication products.
- Maintains high data integrity.

About Diligent, an IBM company

About Diligent Technologies:

Diligent was acquired by IBM Corporation in April 2008. Diligent delivers scalable, proven, enterprise-class disk-based data protection solutions. Diligent's ProtecTIER™ platform, powered by HyperFactor™, a game-changing inline data de-duplication technology, enables customers to protect more data while recording less of it, all without disrupting existing policies, practices and procedures in their data center. For more information about Diligent, please visit http://www.diligent.com

About IBM:

For more information about IBM System Storage, please visit http://www.ibm.com/storage

©2008 Enterprise Management Associates, Inc. All Rights Reserved.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst and consulting firm dedicated to the IT management market. The firm provides IT vendors and enterprise IT professionals with objective insight into the real-world business value of long-established and emerging technologies, ranging from security, storage and IT Service Management (ITSM) to the Configuration Management Database (CMDB), virtualization and service-oriented architecture (SOA). Even with its rapid growth, EMA has never lost sight of the client, and continues to offer personalized support and convenient access to its analysts. For more information on the firm's extensive library of IT management research, free online IT Management Solutions Center and IT consulting offerings, visit www.enterprisemanagement.com.

This report in whole or in part may not be duplicated, reproduced, stored in a retrieval system or retransmitted without prior written permission of Enterprise Management Associates, Inc. All opinions and estimates herein constitute our judgement as of this date and are subject to change without notice. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©2008 Enterprise Management Associates, Inc. All Rights Reserved. EMA™, ENTERPRISE MANAGEMENT ASSOCIATES®, and the mobius symbol are registered trademarks or common-law trademarks of Enterprise Management Associates, Inc.

Corporate Headquarters:
5777 Central Avenue, Suite 105
Boulder, CO 80301
Phone: +1 303.543.9500
Fax: +1 303.543.7687
www.enterprisemanagement.com

1731.110408