A WHITE PAPER BY ITCENTRIX

Hitachi's Data De-duplication Appliances

Extending Performance and Scalability to the Midrange Market

Executive Summary

It was nearly eighteen months ago that ITCentrix first wrote about Hitachi's first introduction of Virtual Tape Library (VTL) solutions with Diligent's ProtecTIER™ data de-duplication technology. We said then it was a "wow" technology. Since then Hitachi has been very successful in establishing itself as a leader in enterprise data de-duplication VTL solutions.

Hitachi's original de-duplication VTL solutions were individually configured for each of its customers. With a maturing market, Hitachi has recently introduced three VTL appliances for the Enterprise, Large Business and Medium-sized Business markets. We believe these new offerings will reduce the cost of support for customers and allow for quicker deployment.

While Hitachi has been very successful in large organizations, smaller hash-based systems from a variety of vendors have dominated small and medium-sized organizations. Hitachi's announcement of a Midrange VTL appliance is especially interesting, as it competes with the hash-based systems on price but provides much greater scalability and flexibility for smaller and midsized data centers.

De-duplication has evolved from an esoteric to a "must-have" technology in less than eighteen months.

With the explosive growth of unstructured data, most environments have experienced enormous growth in data that must be protected. Backup and recovery remain the number one problem in the minds of storage managers. This paper concludes that in-line data de-duplication systems in general and Hitachi's ProtecTIER-based VTL appliances specifically can significantly reduce costs and improve recoverability in most data centers.

This paper analyzes the differences between two in-line de-duplication architectures, hash-based and HyperFactor (the technology inside ProtecTIER). It concludes that for smaller data centers, both architectures work well. For organizations whose backup requirements are growing significantly and that have ten terabytes or more to back up, ITCentrix research shows that the economics of using ProtecTIER are more attractive. It is very likely to require fewer VTL units to deploy and manage, less storage and lower operational costs than either purely tape-based systems or hash-based de-duplication VTLs.

As an example, this paper includes a case study in which the Hitachi solution gave 31% greater savings than an equivalent hash-based VTL system in a midsized environment. In this example, the customer saw nearly $1 million in net present value (NPV) and an

internal rate of return (IRR) of nearly 100%. Perhaps as important, the customer feels its environment is much more flexible to meet change and is far better positioned to accommodate the explosive growth in data that needs to be protected via disk-based backup.

Data de-duplication has evolved from an esoteric to a “must-have” technology in less than eighteen months. The must-have, as always, has to be directed at solving the right business problem. While ProtecTIER is not the best fit for all environments, ITCentrix strongly recommends that IT executives responsible for fast-growing data requirements and with more than ten terabytes to back up regularly should evaluate the technology as a fundamental component of their backup and recovery strategy.

ITCentrix
February 2008

TABLE OF CONTENTS

Hitachi Product Announcements

Hitachi's Virtual Tape Libraries use ProtecTIER™ data de-duplication technology from Diligent Technologies. This technology is the market leader in the enterprise space, and Hitachi has been very successful in marketing the solution to enterprise customers. Until now, Hitachi custom-designed solutions for its customers using its own storage, multi-core servers and ProtecTIER software. Hitachi has announced three appliances that simplify its go-to-market approach, particularly in midrange markets. The appliances are integrated solutions consisting of a server, operating system, storage and data de-duplication software that can be ordered as a single part number. The three appliances are:

…these appliances …make the solution quicker to configure, order and implement, and easier for the sales force and channel partners to administer.

1. Virtual Tape Library Appliance for Medium Business
    - AMS 500 with 8TB of FC & SATA Drives
    - Expandable to 20TB
    - ProtecTIER Midrange Software
    - Multi-core Server
    - Target Performance = 200MB/sec
2. Virtual Tape Library Appliance for Large Business
    - AMS 1000 with 15TB of FC Drives
    - Expandable to 30TB
    - ProtecTIER Enterprise Software
    - Quad Dual-core Server
    - Target Performance = 300MB–400MB/sec
3. Virtual Tape Library Appliance for the Enterprise
    - AMS 1000 with 30TB of FC Drives
    - Expandable to 50TB
    - ProtecTIER Enterprise Software
    - Quad Dual-core Server
    - Target Performance = 400MB/sec

The advantage of these appliances compared with their predecessors is that the solution is quicker and easier for sales and channel partners to configure and order and are user installable.

Hitachi will continue to offer custom configurations to meet the specifications of customers with special requirements. One such popular configuration is a fully populated USP V with four ProtecTIER servers attached.

ProtecTIER Technology

Virtual Tape Libraries employing data de-duplication are primarily used to improve backup and restore systems, because essentially the same data is backed up many times. The backup software (e.g., NetBackup, TSM, Networker) sends data to the ProtecTIER VTL server as it normally would to a standard automated tape library. The server holds a memory-resident index of data previously stored in the array. The HyperFactor de-duplication algorithm determines which data previously stored is most

similar to the data currently being backed up. The previously stored data is read back from the array and compared bit-for-bit with the new data being backed up. If the data being backed up is unique, the data will be stored as normal on the disk array. If the HyperFactor algorithm determines the backup data is a duplicate of data previously stored, it will not write the data into the repository but reference the previously stored data. This backup method ensures that only unique data is stored on disk, thereby minimizing the amount of storage required.

The key factors in determining the amount of data reduction are:

- The degree to which one version of backup data is similar to another (degree of change). The higher the degree of similarity, the greater the data reduction.
- The frequency of backups. The higher the number of backups in a given time period, the greater the data reduction.
- The retention period. The longer the retention period, the greater the data reduction.

In real data center environments, ProtecTIER can be expected to achieve a 10x-to-20x reduction in the amount of data to be stored in the VTL.

In real data center environments, ProtecTIER can be expected to achieve a 10x-to-20x reduction in the amount of disk needed to store backup data in the VTL. Figure 1 shows the basic schema.

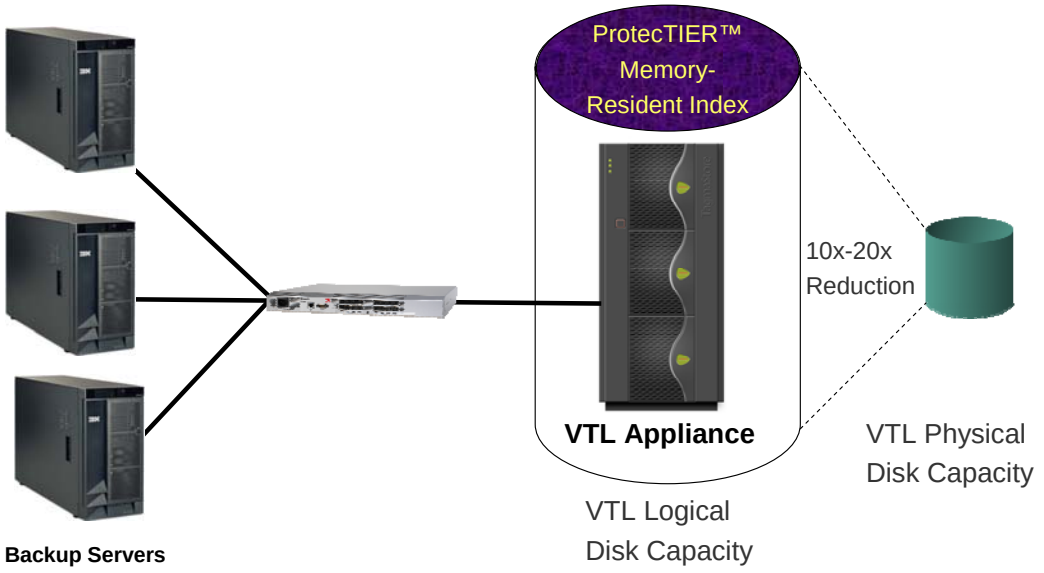

Figure 1 – Hitachi's Virtual Tape Library Appliance with ProtecTIER (*A 10x-to-20x reduction is a good practical range for planning purposes – higher figures [up to 33x] are often quoted by some vendors but are unlikely to be achieved in production environments.)

User discussions with ITCentrix analysts have confirmed that most databases are excellent candidates for data de-duplication. Email systems with long retention periods and DB2/Oracle databases were found to show good reduction ratios. Files such as text files or SQL backup files usually have lower reduction ratios. Compressed and encrypted files result in very low (or negative) reduction ratios.

Comparison of Data De-duplication Technologies

Backup Objectives

There are a number of objectives for backup and restore processes. The principle ones are:

1. Protect the data from a local disaster. If there is a local disaster, some data will be lost (unless expensive synchronous remote replication schemes are used). The quicker the data can be moved off-site, the less data will be lost when the system is restored. Most organizations set a recovery point objective (RPO) for applications. They take into consideration the impact on the business from lost data and the cost of reducing the amount of data lost and derive an optimum RPO.

2. Provide a robust method of restoring data and systems. This may be required because of a local disaster, as mentioned above, or may be due to corruption of a file, a program malfunction, a table accidentally deleted or a myriad of other reasons. The quicker and more reliably the data can be recovered and the system restored, the lower the impact on the business. Organizations set a recovery time objective (RTO) for applications.

Backup systems have to be designed and organized to meet these objectives. In particular, the processes and procedures for a backup system have to be rock solid. The backup and recovery system itself has to be very reliable. Data center managers are rightly very reluctant to change backup and recovery processes that work. Change is risky.

Post Process (aka Off-line) Data De-duplication

A post process data de-duplication system has the advantage that the initial "ingest and protect" phase of a two phase process can be completed faster than an equivalently configured inline de-duplication system, resulting in a better RPO. There is also no change required to the protect phase.

However, a disadvantage of this approach is the introduction of a "de-duplication window" that must be managed in addition to the "backup window" administrators are often struggling to meet. The length of time required to complete both the "ingest/protect" and "de-duplication/recover" phases of this two step operation is always longer than the inline approach. It matters how much time is required to complete backup operations, vault backup data to tape, get them off site, and be in a position to be able to restore data from a disk-based VTL before the start of the next day's backup-cycle. In many backup environments, the two phase post process de-duplication approach will not be able to complete within a 24 hour period without deploying a large number of de-duplication servers. This is typically not acceptable to organizations responsible for protecting mission critical data.

Another disadvantage of the post processing approach is that it takes significantly more disk storage to hold or process the data before it is de-duplicated. Disk contention is another potential problem; disk performance could be reduced as users attempt to access storage during the de-duplication process.

A final and probably most important reason that off-line de-duplication is usually the wrong process model, because the recovery procedures are much more complicated than those with in-line data de-duplication. With an in-line model, the processes are under the complete control of the backup and recovery software. With an off-line model, additional software and processes have to be added to connect (say) the tape library to the VTL and manage that process. There is no longer a single point of control for backup and recovery. This method may be optimum for some types of archiving (when the data is unlikely to be read again). However, for backup and recovery of important mission-

critical applications that have aggressive RPO and RTO objectives, it is rarely appropriate or used.

In-line Data De-duplication

In-line de-duplication removes duplicate data from the backup stream in real time as it is being received from the backup server or backup disk pool. This approach minimizes the amount of physical disk required, since no additional storage space is required to temporarily hold un-de-duplicated data. The example in Figure 1 above shows the backup data being sent by the backup server(s) directly into the ProtecTIER server, where it is de-duplicated in real time before it is stored on disk.

The point in the backup process at which the de-duplication process is inserted is an important design consideration. Where RPO is not the most critical requirement and speed of the overall process is critical, the VTL server can be inserted at the beginning of the process as in Figure 1. However, for organizations that have very aggressive requirements for RPO and/or want to minimize the risk of inserting a new technology into the backup and restore processes, the de-duplication process can be inserted later. For example, in the case study the first step is to produce a Tivoli Storage Manager (TSM) backup pool that is used to create a remote copy of data over the network. The second stage is to put that data into the VTL directly from the TSM pool. The TSM software is in charge during the whole backup and recovery process, including the in-line de-duplication process.

De-duplication Technology Types and Comparisons

De-duplication systems using hash-based architectures max out at about eight to ten terabytes and are suitable for small and midsized data centers.

There are two major types of de-duplication technology. They are systems that deploy hash-based algorithms, and systems that use Diligent's HyperFactor algorithm.

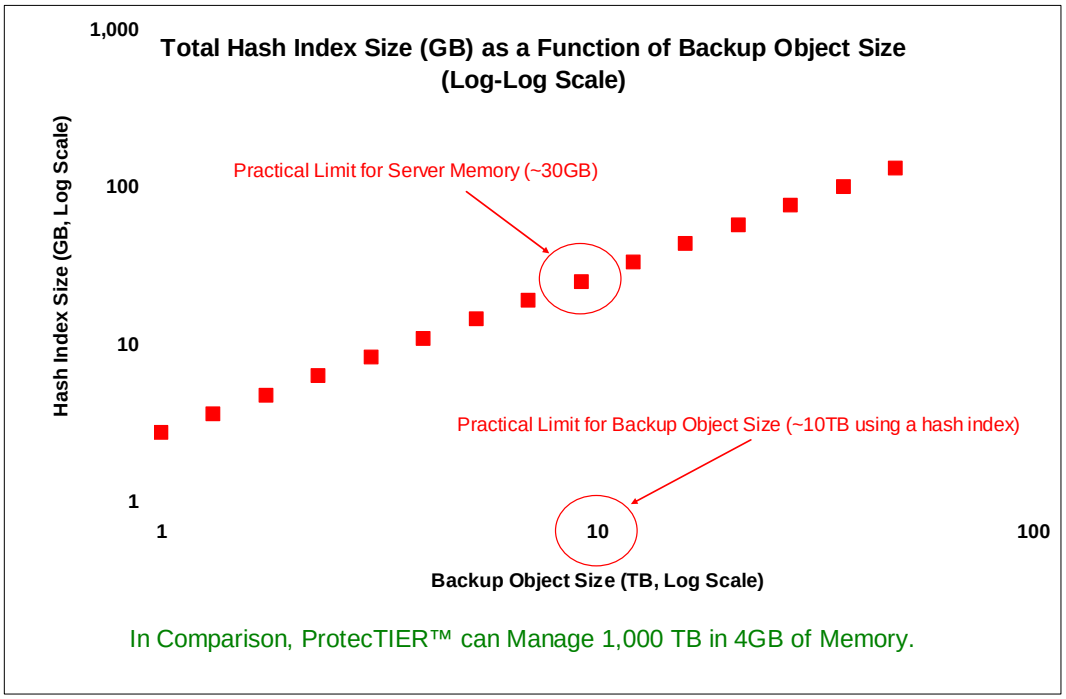

Figure 2 – Total Hash Size

The technical details below demonstrate why Hitachi's Virtual Tape Library solutions using Diligent's HyperFactor algorithm can take advantage of the benefits of in-line data de-duplication even for large amounts of data, whereas competitive VTL systems that

use a hash index approach can generally be used in-line only for small applications and usually design post-process solutions for larger applications or in enterprise data centers. Figure 2 above graphically shows the major differences between a hash-based approach and ProtecTIER's HyperFactor architecture. In the hash-based system, the size of the hash index grows linearly with the amount of data stored in the VTL repository. This approach is acceptable for small amounts of data, as the hash index can reside in server memory. Once the hash table grows to the point where it can no longer be held in RAM memory in the server, the performance of the in-line data de-duplication process dramatically degrades because the index must be paged to disk. Figure 2 shows that hash-based systems cannot deal with backup capacities greater than ten terabytes in size.

Technical Detail

The ProtecTIER HyperFactor technology is a major advance over traditional hashing algorithms in a number of areas.

1. By being able to look at data over any appropriate length, a petabyte of data (1 petabyte = 1,000 terabytes) can be represented within a four-gigabyte memory-resident index. The matching can take place at memory speed with a single pass of the data. This means a single instance of ProtecTIER can scale to support many hundreds of terabytes, is much faster and is an effective solution in the largest enterprise data center.
2. Because new data to be backed up is compared byte for byte with existing data, there is no possibility of any error because of a false positive match and no possibility of data corruption.
3. In contrast, de-duplication systems that use hashing algorithms (e.g., MD-5, SHA-1, SHA-2) split the data up into "chunks." If (say) a ten-terabyte object has to be backed up and the size of the each chunk is (say) 8K bytes, there will be a total of 1.34 billion hash numbers that have to be compared. Each hash is twenty bytes long, so the size of the hash index is 26.8 gigabytes in size. Unless this is held in server memory, the performance slows down enormously. Thirty gigabytes of memory is the maximum size of main server memory that is practically available. De-duplication systems using hash-based architectures are slower and max out at about ten terabytes and are suitable for small and midsized data centers. Figure 2 below shows the impact.
4. Hash-based systems have a very small but finite chance of a hash collision. Although the probability of it happening is very small, the consequences could be very severe. Nassim Nicholas Taleb's book *The Black Swan* illustrates this eloquently. There has to be a plan for this eventually.

To overcome this architectural limit, hash-based architecture vendors have to deploy multiple smaller VTL appliances. The larger the amount of data to be stored, the more appliances must be deployed. The implications of this in the data center are profound. Backups must always be broken into small units and pointed at the same appliance. Data sets have to be broken down into small pieces to fit into the systems, which reduces the potential efficiency of the de-duplication process and increases the amount of storage required. The number of appliances grows very quickly, and each one has to be scheduled carefully. If the systems change in size significantly, the backup of the system must be split across more appliances and the systems be rescheduled, resulting in less efficient data reduction.

ProtecTIER's HyperFactor de-duplication algorithm requires no more than four gigabytes of memory even when the repository consists of one petabyte of physical storage. Ensuring that the HyperFactor index remains in memory at all times results in a ProtecTIER server being able to perform in-line data de-duplication at 400MB/sec regardless of the amount of data stored in the repository, which is over twice the speed of hash-based architectures.

Figure 3 below enumerates the assumptions that are reflected in the comparison charts

Assumptions	ProtecTIER™	Hash-based Architecture
Total Data to Be Protected (TB)	80	80
Hours to Complete Backup	10	10
Percentage of Data Backed Up Each Night (Incrementals and Full)	67%	67%
Large Backup Objects (% Total)	20%	20%
Maximum Backup Object Size (TB) that Can Be Backed Up in Window	14	6
Compression Factor for Small Objects	12	12
Compression Factor for Large Objects	15	13
Data Read/Sec (MB)	400	182
Disk Utilization for 10TB	80%	80%
Disk Utilization for 1,280TB	75%	65%
Staffing/Server	0.1	0.1
Cost/IT Staff/Year	$100,000	$100,000

Figure 3 – Assumptions

throughout this paper. Note that the compression ratios are not the maximum that can be achieved but are reasonable planning figures that can be found across the application types typically found in a data center. Again, the read rates for a de-duplication system are typical rates that can be achieved, not maximum rates that can only be achieved in a benchmark situation. The most important difference in defining the scalability of the systems is the maximum backup object size that can be backed up in the backup window (assumed to be ten hours). The maximum backup object size is fourteen terabytes for ProtecTIER, against only six terabytes for hash-based architectures. This does not matter for small data centers but has a significant impact as the size of data centers increases.

Figure 4 below quantifies the impact of the technology differences on the amount of storage required to hold the de-duplicated data. As the amount of data to be backed up grows, the difference between the two different types of technology increases.

Figure 5 below shows the large difference in the number of de-duplication servers required as the size of installation grows. For large installations, more than twice the number of servers are required if hash-based architectures are used, compared to the number required with ProtecTIER. This means that more people are likely to be required to manage the installation, and storage will be used less efficiently.

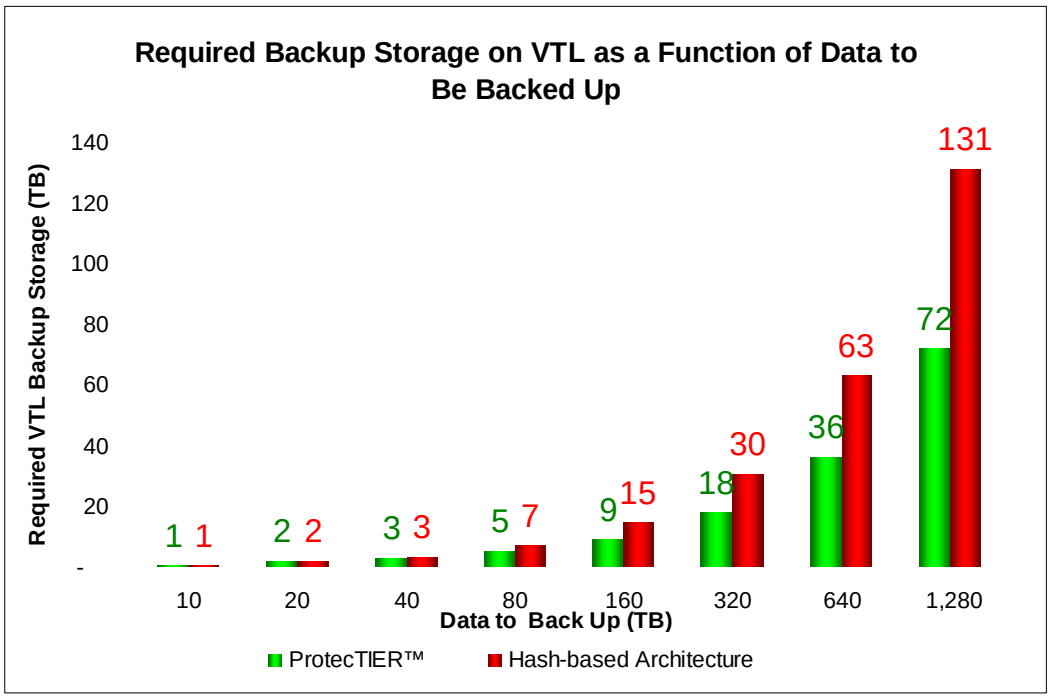

Figure 4 – Storage Required to Backup Data

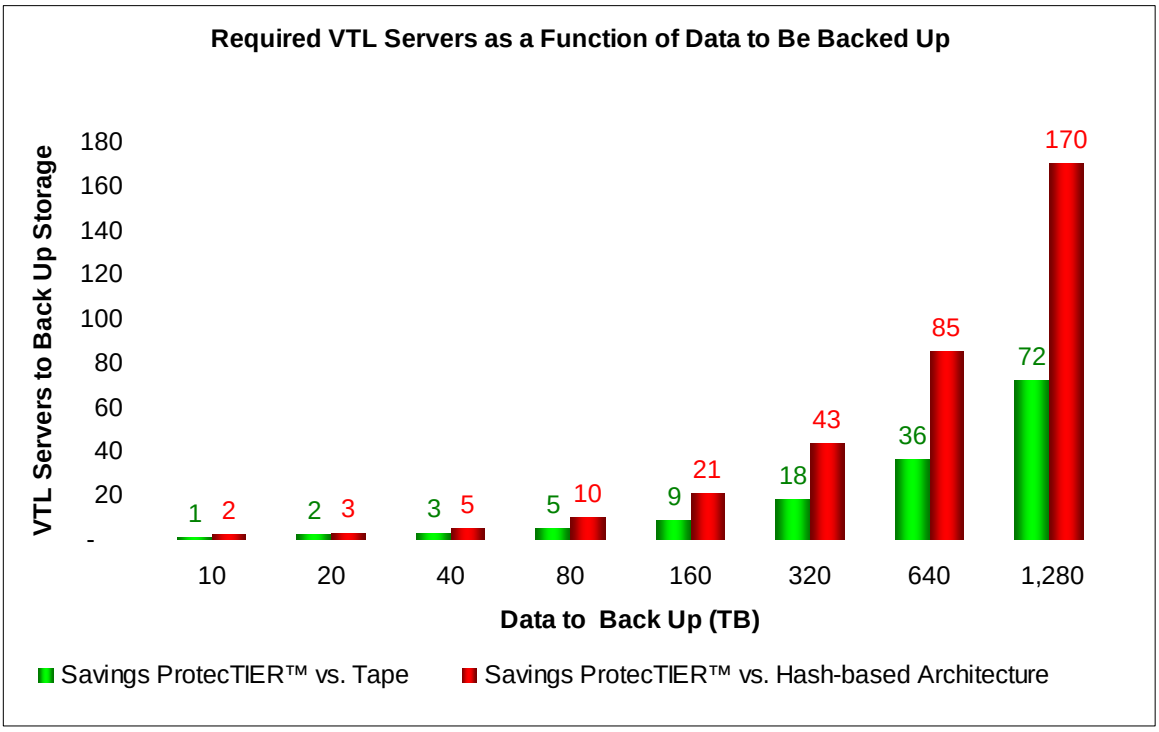

Figure 5 – Number of De-duplication Servers Required

Figure 6 below shows the difference in the number of staff required to manage the VTL de-duplication systems.

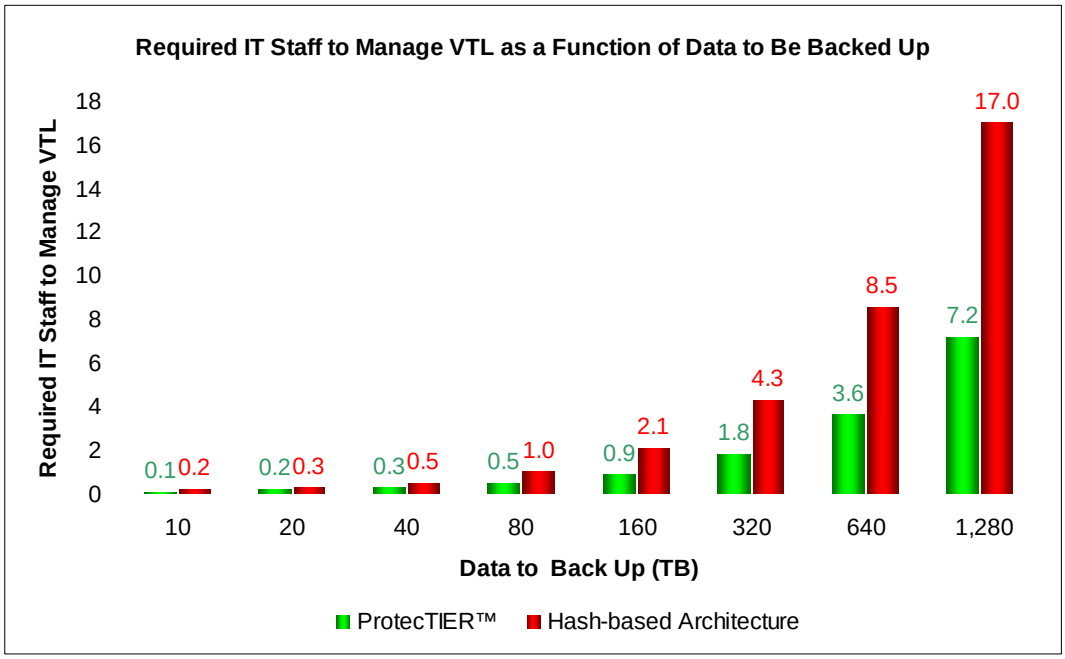

Figure 6 – IT Staff Required to Manage the De-duplication Process

The differences are very modest in small installations but become very significant as the amount of data to be backed up grows.

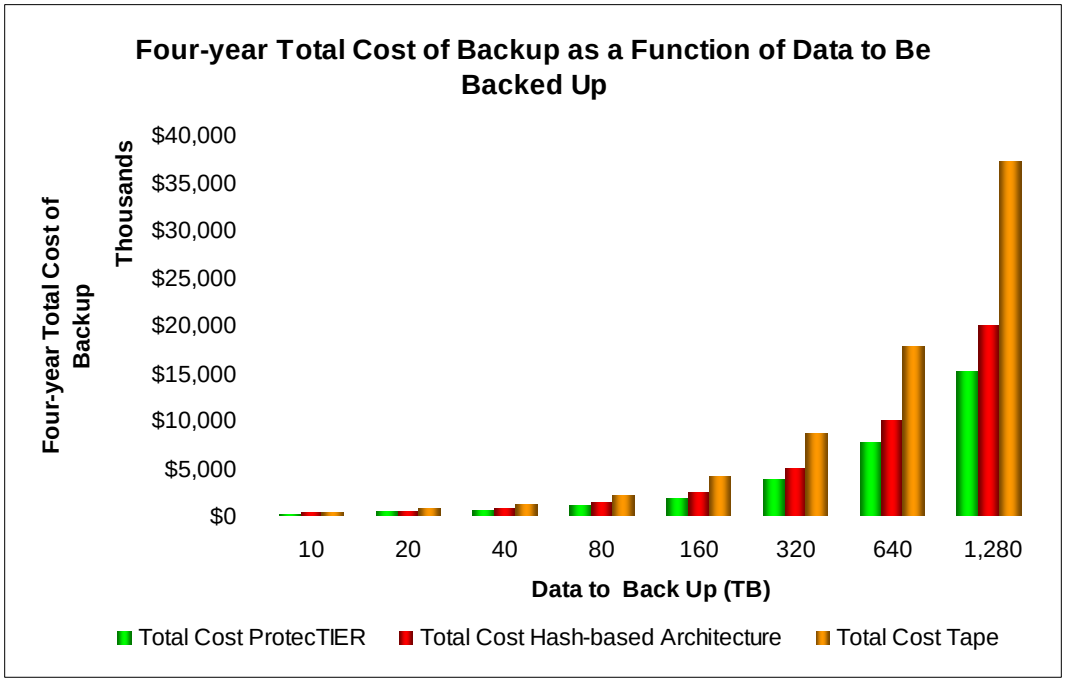

Figure 7 – Total Four-year Cost of Backup

For small amounts of data, both hash-based and ProtecTIER systems show significant benefits over traditional VTL approaches. As the amount of data to be backed up scales, the difference between the ProtecTIER and the other systems grows very significantly.

Figure 7 above shows the total four-year costs of backup using a VTL tape-based system, a VTL with ProtecTIER technology or a VTL with hash-based technology.

For small amounts of data, both hash-based and ProtecTIER's HyperFactor in-line de-duplication show significant benefits over traditional tape and VTL approaches. As the amount of data to be backed up scales, the difference between the ProtecTIER and the other systems grows very significantly.

Case Study

Challenges with Current System

A large US-based financial company was reviewing backup and recovery procedures for the critical systems in its data centers. There were a number of areas for improvement, including:

- Reducing the number of tapes being transferred out of the data center, because of requirements that all tapes be encrypted when in transit
- Retaining or improving the recovery point objective (RPO) for a disaster in the data center
- Significantly improving the ability to restore data for normal everyday problems such as data corruption or tables being accidentally deleted (the time taken to restore systems from the tape libraries was too slow and unreliable and did not meet the organization's recovery time objectives (RTO))
- Improving the cost effectiveness of the backup processes

The total amount of data in these critical systems was approximately eighty terabytes, spread over two primary data centers. The backup software being used was IBM's Tivoli Storage Manager (TSM) with tape libraries, and the disk storage was a mixture of storage arrays from different vendors, with the majority coming from Hitachi.

Solutions Evaluated

The IT task force evaluated a number of disk-based VTL solutions and tape library systems.

Initially, the group focused on methods of meeting the recovery point objectives for the organization. The team evaluated alternative ways of getting coherent sets of data off-site. It was determined that the current tape-based method had to be changed because of the lack of data encryption. One approach evaluated was to de-duplicate the data and then transmit it electronically to a remote site. This approach was very cost effective from an IT perspective, because the cost of transmitting the data was reduced by a factor of ten or more. However, this approach was rejected by the business, because the total elapsed time to de-duplicate the data and transmit it was far greater than that of transmitting the data directly from the TSM backup pools to a remote site. The business would not accept lengthened RPO objectives for these critical applications – quite the opposite; the business is considering improving the RPO significantly with asynchronous remote replication.

The second focus area for improvement was RTO, which involved improving the recovery of data and rerunning or restarting systems. Senior IT management mandated that there was to be no major change to the fundamental backup processes and procedures; TSM was to continue to be the system supporting backup and restore. The task force quickly

narrowed the technology options to be investigated to virtual tape libraries with de-duplication functionality. They narrowed the investigation still further to compare two in-line de-duplication solutions against their current system:

- In-line VTL de-duplication solution using ProtecTIER
- In-line VTL de-duplication solution using a hash-based architecture
- Current tape library system upgraded

The team quickly concluded that disk-based systems offered a significantly better basis for recovering data faster and more reliably, which eliminated upgrading their tape library. The addition of de-duplication to VTLs meant that far more data could be held online to be available for recovery purposes. The only decision left was what type of de-duplication technology should be used.

Solution Specification

The solution developed was to initially back up the critical application data to TSM disk pools and then transmit this data to a tape library at the remote site. Each site backed up to the other site. This eliminated the need to produce tapes and transmit them physically and met the business RPO requirements. This was to be completed within the initial ten-hour backup window.

The data in the TSM disk pools is available to be used for recovery purposes. To improve recovery time objectives, the data has to be held longer, at a lower price point. To achieve this, the data is transferred to long-term disk-based VTL from the TSM disk pool. De-duplication is done to increase the amount of data that can economically be held and therefore allow much faster and more granular ability to recover a file or a system.

The team evaluated two VTL solutions, a ProtecTIER-based solution from Hitachi and a leading hash-based system. This part of the process had to be completed within a second ten-hour window, in order to be able to release the TSM disk pool in time to be used for the next day's backup.

VTL Comparison

The team compared in detail the performance, throughput and ease of use of the solutions based on the two architectures, as shown in Table 1 below:

Metric	ProtecTIER	Hash-based Architecture
Throughput (MB/sec)	400	182
Compression Ratio (small DB)	12	12
Compression Ratio (large DB)	15	13
# Servers Required	7	16
VTL Terabytes to Store 80TB	8.2	10.7
# Operational Staff Required	0.5	1.0

Table 1 – Comparison of VTL Architecture to Back Up 80TB

The throughput of the ProtecTIER solution was significantly better than that of the hash-based architecture. In addition, it could handle much larger backup objects, which allowed greater compression ratios for large objects to be backed up with lower storage costs. The ProtecTIER solution allowed much greater operational flexibility and fewer points of control, which resulted in lower operational staff requirements.

Business Case

The team compared the configurations of the two architectures in detail, as shown in Table 2 below. The hardware and software costs were assumed to be the same – the difference was the reduced amount of storage required, priced at $20/GB. The maintenance was assumed to be 18% of the hardware and software prices. The staffing requirement was based on the number of servers required and the operational complexity of scheduling the workload against them.

The benefits of the system include reduced tape handling and increased speed and reliability of data recovery (improved RTO). These benefits were assessed by a different group and were assumed to be the same for both architectures.

Financial Analysis	ProtecTIER	Hash-based Architecture
Server Hardware & Software	$1,000,000	$1,033,523
Maintenance (4 years)	$720,000	$744,136
Staffing ($100,000/year)	$200,000	$400,000
Total Costs	$1,920,000	$2,177,659
Total Benefits	$3,022,400	$3,022,400
Net Benefits	$1,102,400	$844,741
% Additional Benefit	31%	
Net Present Value (NPV)	$780,318	$556,996
Internal Rate of Return (IRR)	96%	65%
Breakeven (Months)	23	27

Table 2 – Financial Analysis of Alternative Architectures

Table 2 shows that both de-duplication VTLs offer a good financial case against keeping the current tape-based system. The ProtecTIER solution gives 31% additional benefits over the hash-based architecture, with an NPV of $1.1 million, an IRR of 96% and a breakeven of just under two years.

Conclusions from the Business Case Analysis

The simple conclusion from the analysis is that ProtecTIER was a far superior, more flexible and scalable solution, with the ability to be extended to other systems within the data center. The evaluation team recommended the ProtecTIER solution, and the recommendation was accepted. The system has been in production for about a year and has performed to specification.

The IT department is looking at a number of ways to improve the backup system, including the options of asynchronous remote replication of the primary disk storage for critical systems. This would dramatically improve the RPO and eliminate the transfer over the network. This would then allow the backup process to go straight to the Hitachi ProtecTIER VTLs, which would lower the cost and improve the responsiveness of data and system recovery.

Conclusions and Recommendations

Virtual Tape Library systems with data de-duplication technology are now well established in the marketplace. In the past, the benefits of data de-duplication have been oversold by many vendors, especially concerning the degree to which they can eliminate all tape. However, there is a solid business case for using data de-duplication as part of a disk-based backup and restore strategy.

For backup systems, in-line data de-duplication is effectively the only technology architecture that can meet the elapsed time requirement of most enterprise data centers. Within in-line de-duplication systems there are two major architectures, hash-based and the ProtecTIER approach.

For smaller shops and distributed data centers either technology will work well. However, for large data centers and for fast-growing midsized data centers with more than ten terabytes of storage to be protected, the ProtecTIER architecture demonstrates significantly better economics. It scales better, requires less systems, will usually have better compression ratios and will typically require fewer people to manage.

ITCentrix recommends that IT management forecast the requirements for data de-duplication virtual tape library systems over a three-to-five-year period when evaluating technologies for enterprise backup systems and consider all the costs associated with backup. While the allure of disk-based backup in and of itself will improve total cost of ownership substantially, ProtecTIER can provide an additional layer of meaningful savings and incremental business value for organizations.

ITCentrix
4 Mount Royal Ave., Suite 220
Marlborough, MA 01752
Tel: 774-463-3400
www.itcentrix.com