HP StorageWorks D2D4000 Backup System

A report and full performance test on Hewlett Packard's SME data deduplication appliance

Testing conducted and report compiled by

Binary Testing Ltd Unit 33 Newhaven Enterprise Centre
Denton Island Newhaven East Sussex BN9 9BA

T + 44 (0)1273 615270 E info@binarytesting.com

Introduction

Data storage demands are being driven at a tremendous rate across a wide range of businesses making it essential for administrators to manage existing resources more efficiently. Simply throwing more storage hardware at the problem is no longer an option as businesses must strictly control IT expenditure to remain competitive. There are a number of solutions to this dilemma but this year the hottest of them all is data deduplication.

This technology makes more efficient use of storage facilities by eliminating multiple copies of the same data and reducing them to a single instance. The majority of business IT environments will be providing a range of services to their users including file sharing, messaging and database access and in these scenarios the potential for duplication of data is high.

Workgroups collaborating on a project will have their own versions of files stored which may well be identical or very similar to those maintained by other co-workers. Emails sent to multiple recipients will contain the same message content and file attachments as well.

Take this down to the block level and the opportunities for removing redundant copies of data increase exponentially. When editing documents it's only natural for users to want to keep preceding versions as a safety precaution. Multiple copies of a file will be retained where each one may be only slightly different to its predecessor.

Block level data deduplication technology looks particularly appealing when applied to data backup. Data protection strategies containing a mixture of full and either incremental or differential backups will result in a huge amount of redundant data. Compliance with laid down data protection standards and increased retention periods plus essential secure off-site storage demands only serve to increase the potential for wasted storage resources.

The availability of inexpensive, high-capacity SATA drives means disk based backup is now a viable solution for storing first line backups. However, the philosophy behind tape backup strategies is still being adhered to due to ease of use and reduced management overheads. All of today's backup software products have the ability to fully manage backup and restoration strategies so it makes sense to continue to use these procedures when securing data to disk arrays.

The latest disk arrays add a new dimension to disk based backup by offering up their internal storage as virtual tape libraries (VTLs). Presented as iSCSI (Internet SCSI) or fibre channel targets the server hosting the backup software merely sees them as physical devices locally attached. Backup tasks continue to secure data to cartridges in the library unaware that they are actually virtual network devices. Consequently, as there is no requirement to change existing backup strategies the disk array can be implemented with the minimum of disruption to normal services.

In this scenario there is a clearly defined opportunity for data deduplication. VTL arrays that implement this technology can remove redundant data during the backup phase allowing them to maximise storage potential and reduce costs. Where once large capacity arrays were needed to store first line backups these requirements can be reduced significantly resulting in far lower operating costs, easier management and much more efficient usage of storage media.

Executive Summary

In a recent report the market research company IDC stated that storage capacity is exploding at a rate of almost 60% per year and this accelerated growth is forcing IT executives to rethink what type of storage system is best suited for their data. IDC went further with its top ten industry predictions of technology that will arrive in 2008 and along with virtual tape libraries (VTLs), deduplication is identified clearly as being in demand.

In light of these findings Hewlett Packard's new StorageWorks D2D4000 Backup System is a timely addition to its storage portfolio as it amalgamates both virtual tape libraries and data deduplication into a single solution. Aimed at small to medium enterprises (SMEs), remote offices and businesses with small data centres, the D2D4000 appliance is capable of presenting multiple VTL targets over iSCSI and fibre channel.

For this product HP has opted for in-line, or dynamic, data deduplication as it allows it to provide a good level of features and yet keep a keen eye on costs as well. By using this method the appliance can support any data type and backup application so it can be integrated into existing data protection strategies with the minimum of disruption.

The aim of this report is to take an in-depth look at the StorageWorks D2D4000 Backup System in order to determine its suitability for the target market. It will look at installation, deployment, configuration, ease of use, and features. It will conduct a series of real world performance tests to determine the efficiency of the appliance's deduplication algorithms and to discover whether the claimed ratios of between 20:1 and 50:1 are achievable.

The performance tests will run real world simulations on three popular applications run by the majority of SMEs – file server services, Exchange Server and SQL Server. For each application it will introduce five sets of controlled rates of data change, run automated backups covering a simulated period of three months and will provide a detailed report on the deduplication ratios being returned and the savings made in reduced storage capacities.

This report demonstrates the efficiency of HP's deduplication technology and shows the positive impact it can have on data backup operations. The tests show clearly that the claimed ratios are achievable in the real world with the results ranging from 21.3:1 up to 69.2:1. Storage is also shown to be used far more efficiently with the tests reporting significant disk capacity savings. Using a 4GB sample of data the tests showed savings of between 51.8GB and 229.5GB. Scale this up to a full production environment and it is clear the costs benefits will be considerable.

Deduplication ratios after three month test

Tests	File Server	SQL Server	Exchange Server
0.4% change in 100% of data	**69.2**	**24.2**	**24.9**
1% change in 60% of data	**38.1**	**22.7**	**31.5**
2% change in 40% of data	**31.7**	**22.1**	**37.7**
5% change in 20% of data	**25.5**	**21.3**	**40.3**
10% change in 10% of data	**22.2**	**21.5**	**29.6**

The HP StorageWorks D2D4000 Backup System

HP is currently offering two solutions for the SME market with the family comprising the D2D2500 and the D2D4000 featured in this report. The D2D4000 is a 2U rack mount appliance with room at the front for up to twelve drives with the 4.5TB model under test sporting six 750GB SATA hard disks mounted in the same sturdy hot-swap carriers as used across a wide range of HP's servers and storage appliances.

The D2D4000 is also available with all drive bays populated for a top raw capacity of 9TB. You can choose between dual Gigabit Ethernet ports for iSCSI operations and dual 4Gbps fibre channel ports for integration into an FC SAN – the appliance supports both of these simultaneously. The appliances are shipped with the drives configured in a RAID-6 array which provides dual drive redundancy. There is a higher overhead in terms of storage costs but these arrays can survive the simultaneous loss of two hard disks making for superior fault tolerance. The appliance also comes with both hot-plug power supplies included.

In either IP or FC SAN configurations the appliance can present up to sixteen VTL targets allowing it to support simultaneous backup operations for the same number of servers. Host systems require a standard iSCSI initiator loaded and HP currently supports Windows, Linux, HP-UX 11.23/11.31 and Solaris 10. When configuring VTLs, the D2D4000 offers a good range of options as you can choose from four different library emulations and opt for LTO-2, LTO-3 or LTO-4 tape drives. Each VTL can have up to 24 cartridge slots and also two drives.

Off-site storage is a critical requirement of any backup strategy and HP has this covered as it offers optional kits comprising either LTO-2, LTO-3 or LTO-4 external tape drives and SCSI or SAS host bus adapter cards. For a better fit in your rack cabinet you can also opt for 1U or 3U rack mount kits with the same choice of internal LTO drives. Another alternative is to attach tape drives or libraries directly to your backup servers and allow the backup application to manage securing data to tape.

HARDWARE SPECIFICATION:

Chassis: 2U Rack Mount	CPU: 2 x 3GHz dual-core AMD Opteron 2222
Memory: 8GB 667MHz DDR2	Storage: 6 x 750GB SATA hard disks in RAID-6 array
RAID: HP Smart Array P400 SAS with 512MB cache and BBU	Network: 2 x Gigabit Ethernet
Power: 2 x 750W hot-plug supplies	Management: Web browser

What is data deduplication?

Data deduplication is the process of eliminating multiple copies of data to reduce storage requirements. Deduplication at the file level has a limited appeal but applying this technology at the block level allows redundancy to be reduced further as only unique blocks need to be stored.

There is a wide choice of deduplication technologies available and HP has chosen a particular set to allow it to deliver a cost-effective solution to its target market.

The D2D4000 uses hash-based chunking to eliminate duplication where the appliance examines the data being sent to it and computes a hash value, or fingerprint, for each block of data using an algorithm that uniquely identifies it. These hash values are much smaller than the blocks of data they represent and are stored in an index, or catalogue, file on the appliance.

The hashes are used to compare data being sent to the appliance with that which is already resident. If the hash value for an incoming block matches one in the index then instead of storing it, the reference pointer on the existing block is updated. If the hash value of the block does not match anything in the index then the block is stored on the appliance and its hash value added to the index.

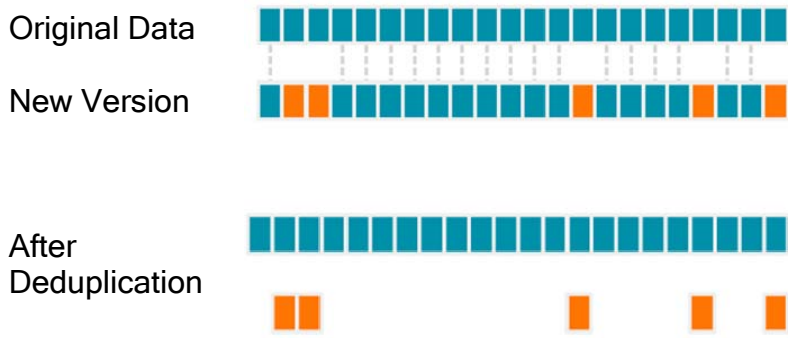

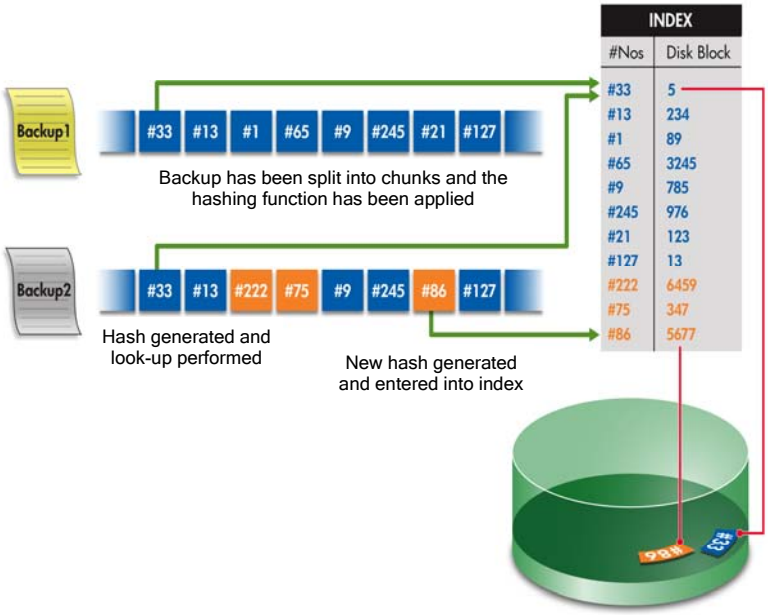

The D2D4000 uses a chunk size of 4k to compute its hash values. To improve performance the index is stored in RAM but even on the D2D4000 with 3TB of usable storage the requirements could be as high as 30GB of memory which is clearly cost prohibitive.

HP has overcome this limitation by introducing a system that optimises memory usage and reduces paging. Hashes are stored in recipe files on the appliance which are compared with each 10MB segment of data being received. It calculates which recipe file it needs to load into memory thus reducing the physical memory requirement substantially.

Deduplication is carried out dynamically at the appliance so no client software is required on the host systems. After deduplication has been applied the appliance then performs compression as it writes the data to disk so delivering even more storage capacity savings.

Dynamic deduplication using hash-based chunking has allowed HP to deliver an appliance-based solution that is physically compact and represents comparatively good value. Furthermore, using a hash index as opposed to methods such as object level differencing allows it to work with any data format.

Installation and Deployment

HP has the initial installation phase very nicely covered as unlike a number of competing solutions it provides a complete wizard-based routine for installing the LTO tape drivers, configuring iSCSI services and discovering the D2D4000 appliance on the network. We tested this process on both Windows Server 2008 and 2003 systems running as domain members and found it could be completed in a matter of minutes.

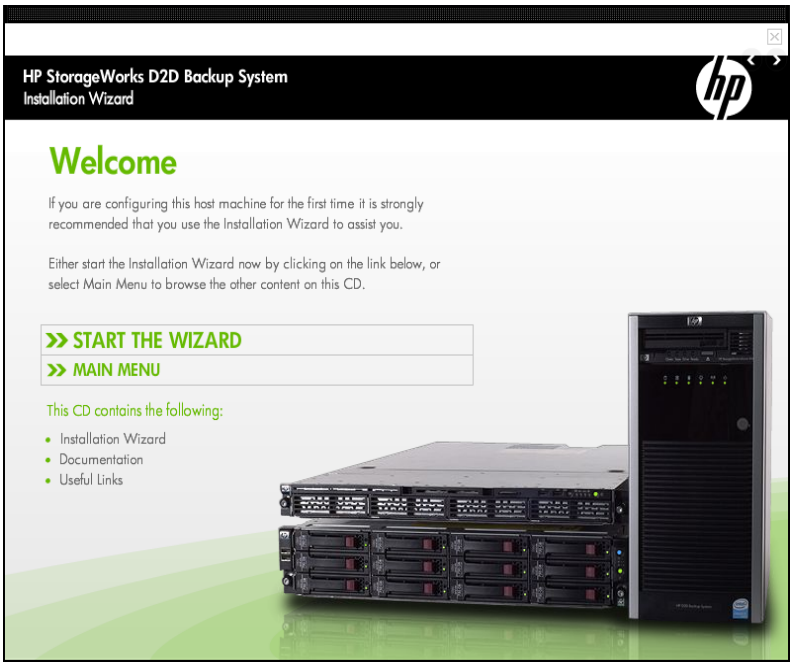

For earlier version of Windows the routine downloads and installs the latest Microsoft iSCSI initiator whilst for 2008 and Vista systems it will simply ask for the resident service to be started by a user with administrative access. The next screen offers options for VTL creation where you can choose from a generic variety or one specifically for Symantec Backup Exec and decide on the number of cartridge slots. If you wish you can also create a standalone LTO tape drive instead. The routine also places an icon on the desktop for quick access to the appliance's web management interface.

Usefully, each time you log an iSCSI initiator on to the appliance it can automatically create a VTL ready for use and assigned to the host's IQN. You can also manually create VTLs and each one may be edited where you can change its name, the type of VTL emulation, the tape drive type and the number of slots. Deduplication can also be enabled or disabled on selected VTLs with a single mouse click.

The home page of the appliance's web interface opens with a complete summary providing an at-a-glance readout on details such network information, general system status and amount of disk space currently allocated to VTL operations.

Hard disk and RAID array status can be monitored and even the condition of the controller's battery backup unit is shown. For each VTL you can view iSCSI target names and aliases for their associated medium changers and tape drives plus details on connected iSCSI initiators.

The main status screen provides plenty of information about the appliance and the Disk tab offers a bar chart showing the total disk space allocated to VTLs, the actual capacity already used and the current deduplication ratio. From the RAID tab you can view installed hard disks, check on their status and activate the beacon function to physically identify them from the front of the appliance.

The Cartridge tab on the management interface shows all the slots for the selected VTL where you can unload all the slots, change their capacities and mark individual cartridges as read only. Bar codes for each cartridge are automatically generated by the appliance and can be viewed and edited from this screen as well.

The Tape Attach option offers a number of features for copying data from the appliance and onto removable media. The Copy operation simply copies data from the selected cartridge slot to a tape loaded in a locally attached tape drive.

Data can be archived with the Export option as this copies data to the tape drive and then deletes it from the VTL. Data from physical cartridges can also be imported to a VTL where a new virtual cartridge will created in a mail slot ready to receive it.

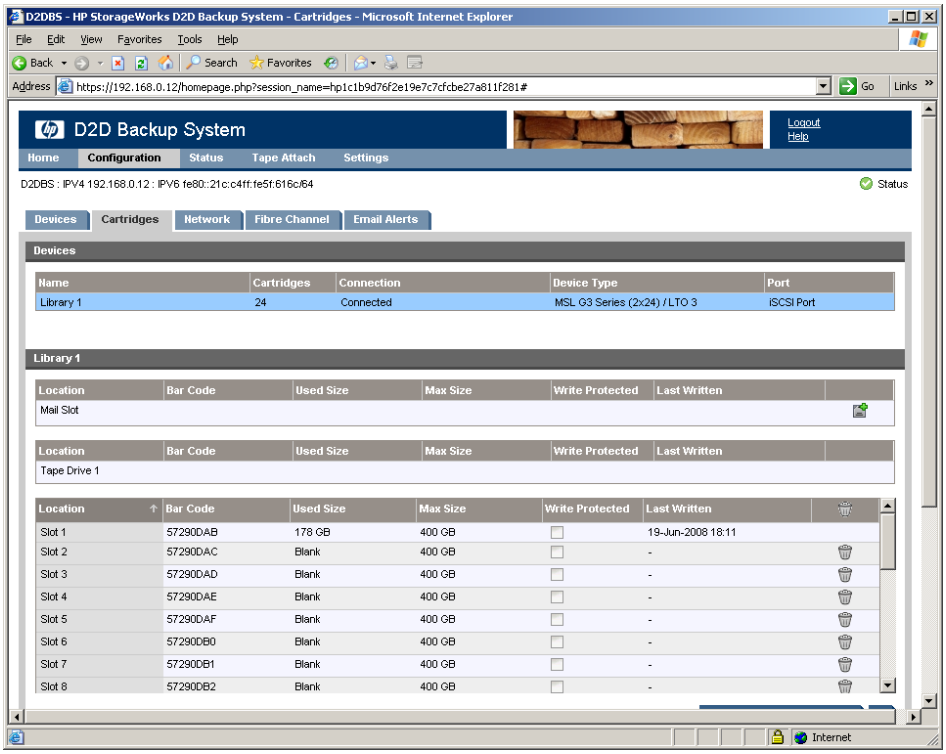

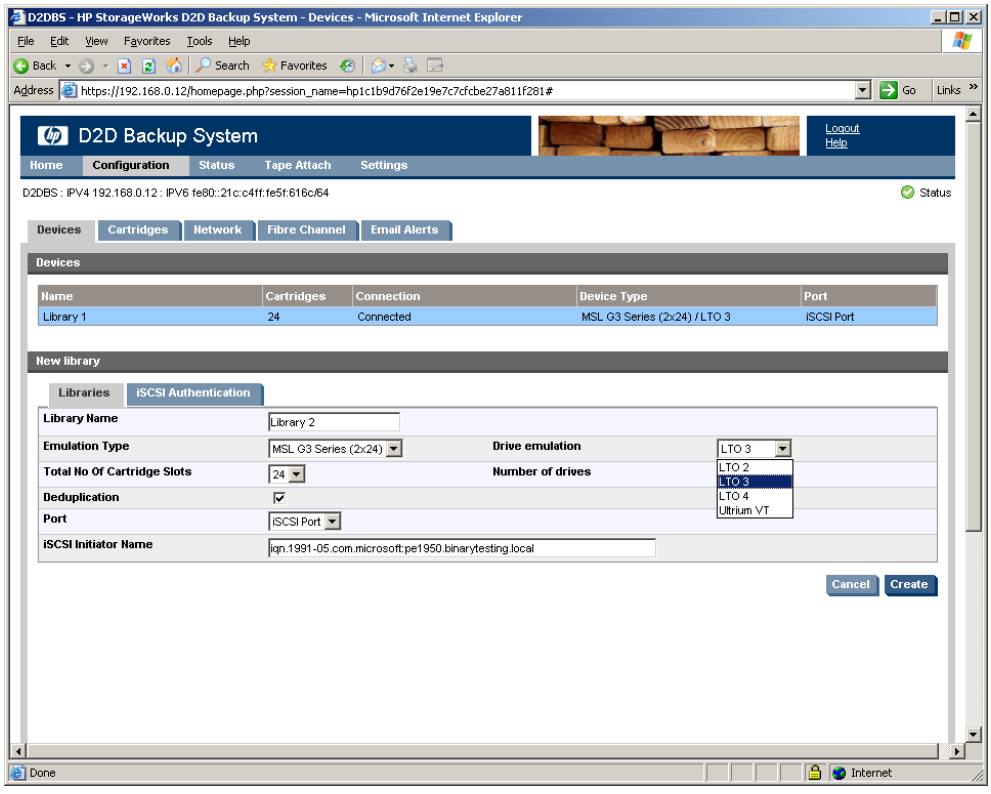

Once our test host initiators were logged in we could see from the Device Manager that a new medium changer and LTO tape drive had been identified and had their drivers loaded successfully. Essentially, the appliance was now ready to go with whichever backup software we wished to use.

We opted for Symantec's Backup Exec v12 which had no problems seeing the test VTL, tape drive and all associated slots. Now all we had to do was create our backup jobs and select the new devices as normal. Overall, we found the entire installation and configuration phase very well handled by HP making it extremely easy to implement the appliance into a backup strategy.

The Test Scenario

Two hardware platforms were selected for the test allowing us to spread the load equally to avoid any contention for system resources. The main domain controller server was configured with Microsoft Windows SBS 2003 R2 and Exchange Server 2003. A second server was configured as a domain member running Windows Server 2003 R2 and Microsoft SQL Server 2005. This system also functioned as the backup server and was installed with Symantec Backup Exec v12 and the SQL Server agent. From here we deployed the Backup Exec Exchange agent to the domain controller.

Test Server Specifications

Domain Controller with Exchange Server 2003	Backup Server with SQL Server 2005
CPU: 1.86GHz quad-core Xeon E5320	CPU: 2 x 1.6GHz quad-core Xeon E5310
Memory: 2GB 667MHz FB-DIMM	Memory: 4GB 667MHz FB-DIMM
Storage: 6 x 750GB SATA II in RAID-5 array	Storage: 2 x 36GB SAS in RAID-1 mirror

To test HP's claimed deduplication ratios of between 20:1 and 50:1 we needed to create a performance testing suite that mirrored real world scenarios closely. We also needed to take into account two key factors that determine deduplication ratios. The rate of change of data and the amount of time it is retained on the storage device will have a profound effect on the ratios.

To this end we devised a group of five tests that allowed us to manipulate the changes to data at a known rate and quantity and apply each of these to a file server, a SQL Server database and an Exchange Server. The tests simulated a three month backup period where the data from all backups for each test sequence was retained on the appliance. One of the advantages of deduplication is its reduced storage requirements allowing data to be retained on first line storage for much longer and the three month period was chosen to show clearly the effects of this.

For the file server test we called for X% of data to be changed in Y% of files. For SQL Server, the Y value would be applied to database rows whilst for Exchange Server this would be applied to attachments sent with individual emails. Five combinations of X and Y values were used to show a range of data changes that could be applied to a number of different real world scenarios.

The file server test data was a 4GB data set held in a number of files. The SQL Server database held 4GB of data in tables, while the Exchange Server had 400 users with message attachments totalling 4GB distributed amongst them.

Rate and quantity of data changes

X% data changed	10	5	2	1	0.4
In Y% of files, database rows or email attachments	10	20	40	60	100

To get any meaningful values for reported deduplication ratios and storage savings we ran each test combination for each application over a simulated three-month period. For the file server test we started by taking a full backup of the directory containing the source data and then followed this by daily incremental backups plus an end-of-week full backup. Once the initial full backup had completed we applied the data change combinations for each day after which the relevant backup task was run.

The backup strategies for each application were selected based on common practices. For SQL Server we ran weekly full backups and daily transaction log backups along with daily differential backups - a best practice defined by Symantec for mid-sized businesses. For Exchange Server we ran daily full backups followed by flushing committed logs.

The Test Scenario

The automated testing software was designed to manage all aspects of the test process - amending the data sets, controlling the type and execution of the backup processes and then extracting the performance data from the D2D4000 device. It was written using a combination of Visual Basic and VBScript. Standard coding interfaces were used throughout to avoid influencing the results by using special low-level coding techniques.

Generating test data can be problematic, and it is difficult to ensure realistic results in terms of content and distribution. We resolved this problem by collecting data from archives of real systems and used this to create the test data sets.

It was important that the tests should not simply change the same piece of data each time, since this would leave most of the data set unchanged and would produce an artificially high deduplication ratio. The data changing algorithms were designed so that each simulated daily run would act on a different part of the data set. This process would continue until either the entire data set used in a test had been altered, in which case the process would simply wrap around to the beginning and continue or the simulated three- month test cycle had concluded.

It was also important that each test's results should not be affected by any previous test run. We ensured this by creating a backup of each data set when we first created it and then used this backup to restore the test data set before beginning each test run.

It was a requirement that the VTL was deleted and a new one created prior to starting each test sequence to ensure a clean slate and also to reset the appliance's reporting facility. Consequently, the data set backups were placed in local storage on the backup server itself. However, we wanted to see how well the appliance dealt with data recovery so conducted restorations of the file server, SQL Server and Exchange Server data sets. These were compared with the local backups and it was confirmed that the data has been successfully restored from the appliance's VTLs in all cases.

Although a test scenario tries to control all the variables, using real-world data and applications makes this difficult in some cases. File systems are generally simple with no hidden overheads but this is not true where databases are concerned, and this extra data can influence the test results.

All database management systems have a data overhead in the shape of control tables and indexes. The database management software uses these to help it to store and retrieve the data and to support SQL, and will change the content as it does so. This control data would be included in a full back up, but not in a differential or transaction log back up.

The management software also carries out its own housekeeping and may choose to reorganise the data on the disk during a test. Re-indexing and block splitting may occur, which will also have an effect on a full back up. The testing software could not control these factors. We were able to reduce the effects of these activities on the results by simplifying the table structures in the database.

We removed all the additional keys from the tables and retained just the primary key in all cases. This reduced the number of indexes in the system and, since we never added or deleted rows in tables but simply updated them the activity in the associated indexes was kept to a minimum. This also reduced the need to reorganise the data, but could not eliminate it completely.

The email tests presented us with the same database issues, but we could do nothing to influence them since we did not intend to manipulate the database directly. However, there was another issue apart from those associated with the Exchange database. Emails cannot be amended once they have been delivered without manipulating the underlying database, and so we would have to follow a different process in order to alter the data in these tests. We therefore produced a test system with two components.

One component simulated a mail server sending messages to the Exchange server using the SMTP protocol, while the other simulated the users logging on to Exchange, reading their mail and then deleting the messages. This component used the POP3 protocol to access the Exchange server.

The Test Scenario

When we set up the test data we created a number of Active Directory users with their email accounts and then sent each of them a number of emails. We then made a full backup of the Exchange database and used this to restore the users and emails at the beginning of each test run. The test then 'read' and deleted the specified percentage of each user's messages, and then sent more emails to replace those deleted, but with amended data based on the percentage change specified.

This had the same effect as amending emails 'in situ' in a similar fashion to that used in the other tests. In this case we used a simple message text for all the emails and added an attachment. The message text remained constant throughout the test, and all data changing took place in the attachment files.

Using the automation scripts, storage usage data on the appliance was collected after every backup had completed and imported into a CSV (comma separated value) file for use in Microsoft Excel. A command line utility provided by HP reported on all storage usage on the appliance and for the performance tests we focused on three key areas.

1) Disk Data Usage = The actual amount of physical storage in MB used to store data from the backup jobs after deduplication has been applied.

2) Cartridge Data Usage = The amount of data being sent to the appliance by the backup task that would be stored if deduplication was not being applied.

3) Deduplication ratio = Cartridge Data Usage/Disk Data Usage.

All graphs were produced using Systat Software Inc. SigmaPlot for Windows 11.0.

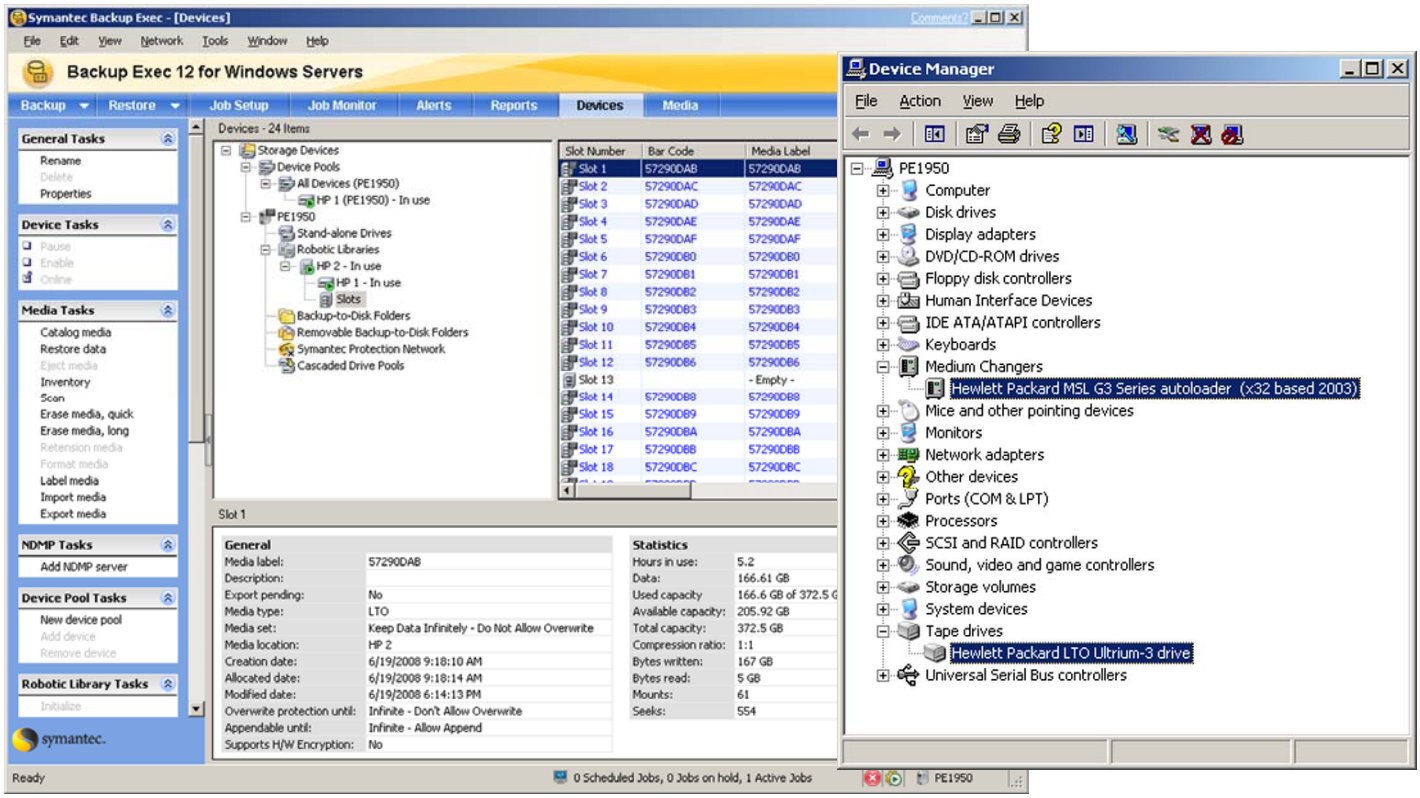

After logging on to the iSCSI targets the host system registered the new medium changer and tape drive in Device Manager and Symantec Backup Exec had no problems working with them during the performance tests.

File Server Test Results

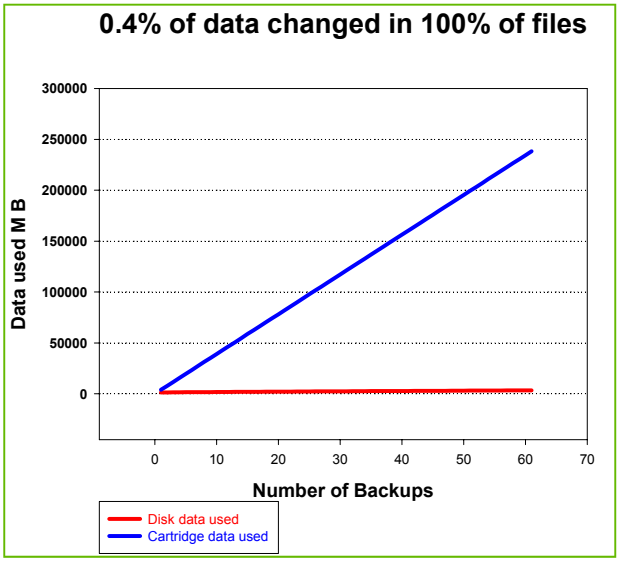

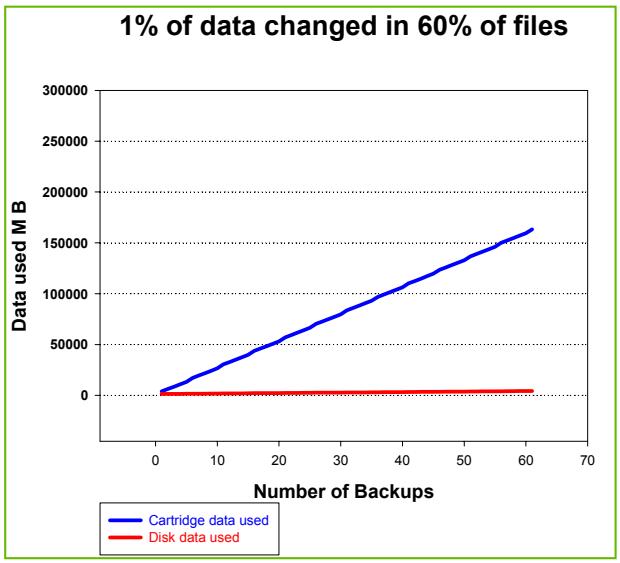

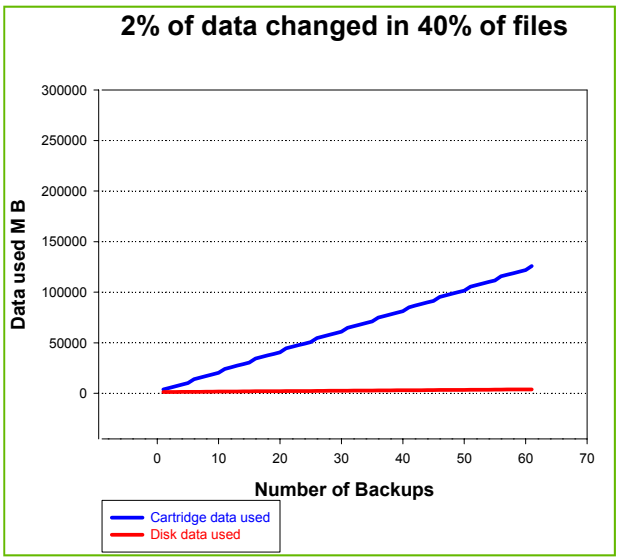

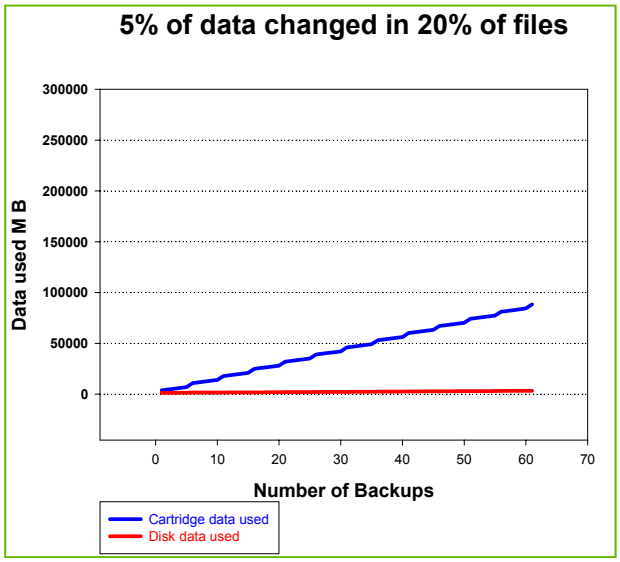

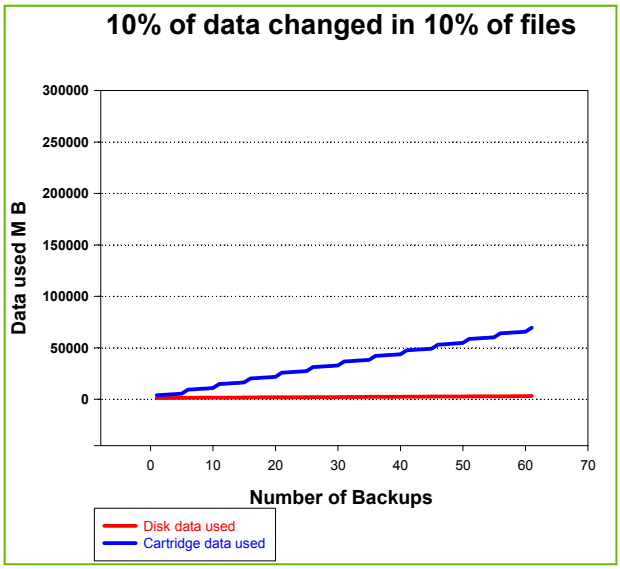

Actual savings in disk storage capacity after 3 months of testing

Test	Disk Data Usage MB	Cartridge Data Usage MB	Savings in MB
0.4% in 100%	3446	238358	234911
1% in 60%	4286	163340	159053
2% in 40%	3972	125831	121858
5% in 20%	3462	88322	84860
10% in 10%	3132	69566	66433

NB: Results for disk data used (red) appear as virtually flat lines on the graphs. This is due to the scaling required to incorporate the much greater values for cartridge data usage in the same graph. The figures shown in the table above confirm this.

Exchange Server Test Results

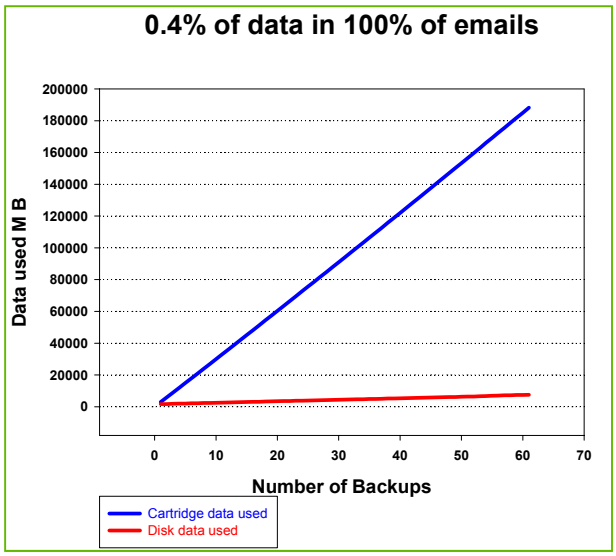

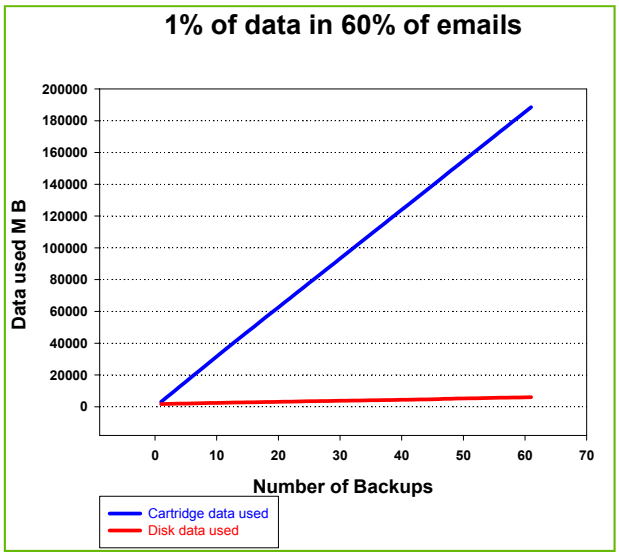

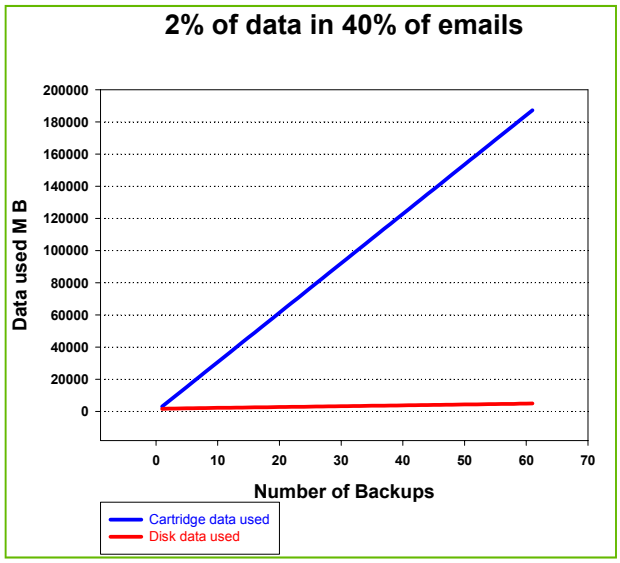

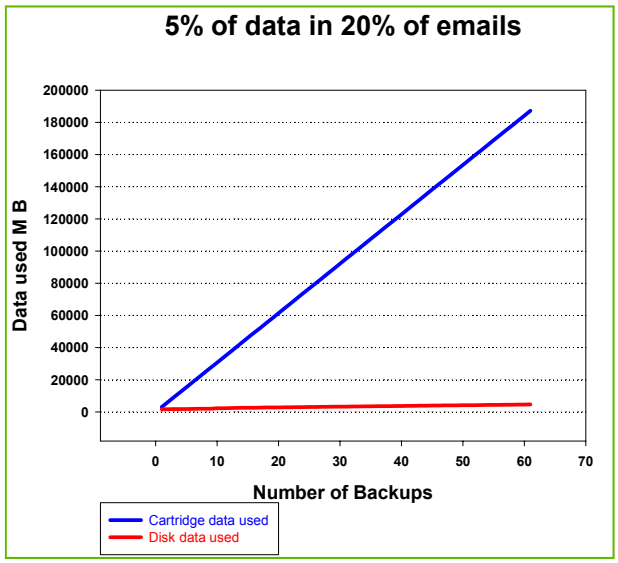

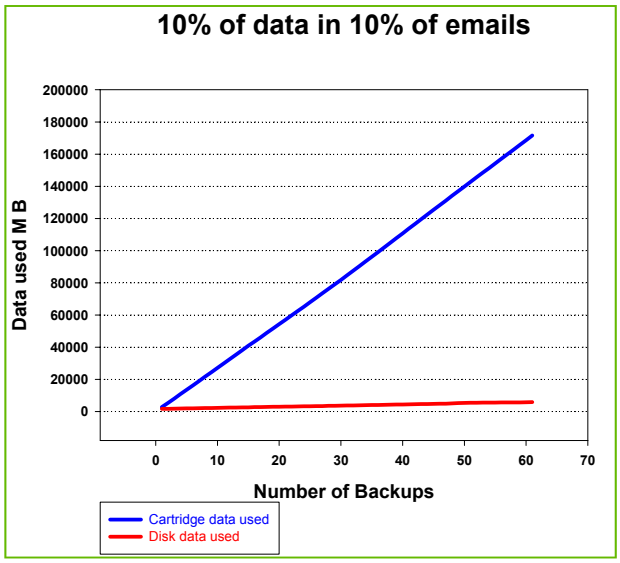

Actual savings in disk storage capacity after 3 months of testing

Test	Disk Data Usage MB	Cartridge Data Usage MB	Savings in MB
0.4% in 100%	7561	188174	180613
1% in 60%	6555	170689	164134
2% in 40%	5552	165354	159802
5% in 20%	6064	172134	166070
10% in 10%	5462	172256	166794

NB: Results for disk data used (red) appear as virtually flat lines on the graphs. This is due to the scaling required to incorporate the much greater values for cartridge data usage in the same graph. The figures shown in the table above confirm this.

SQL Server Test Results

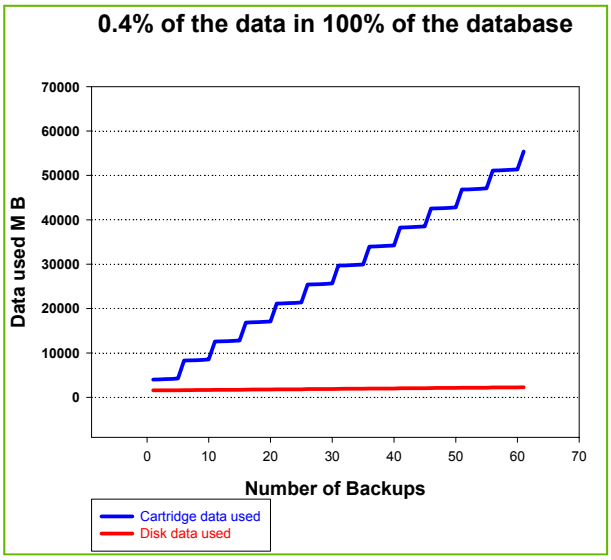

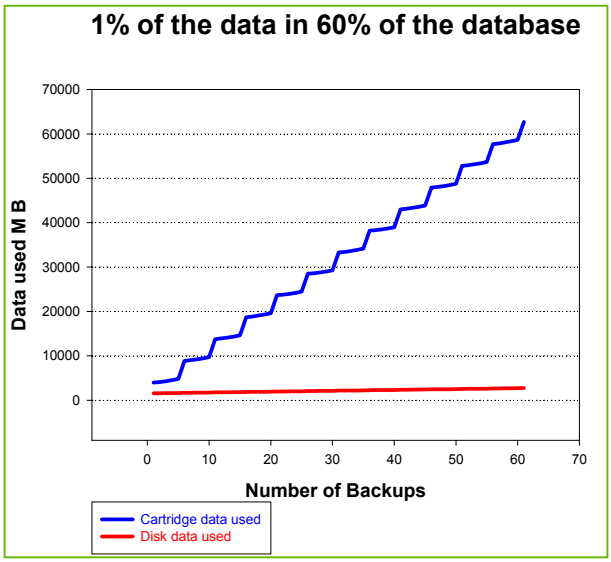

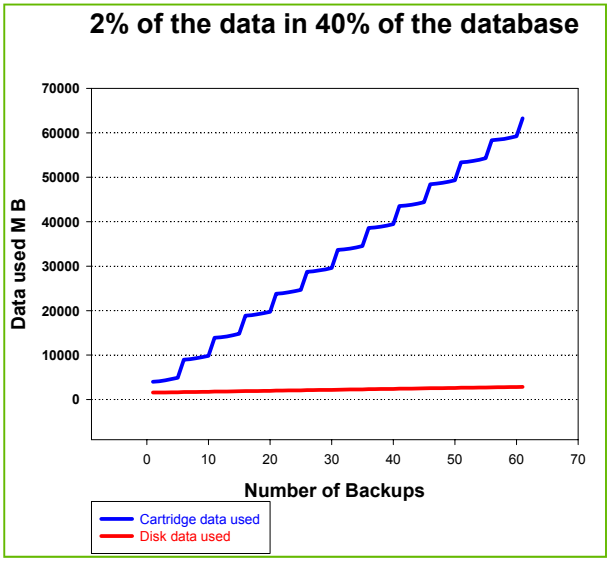

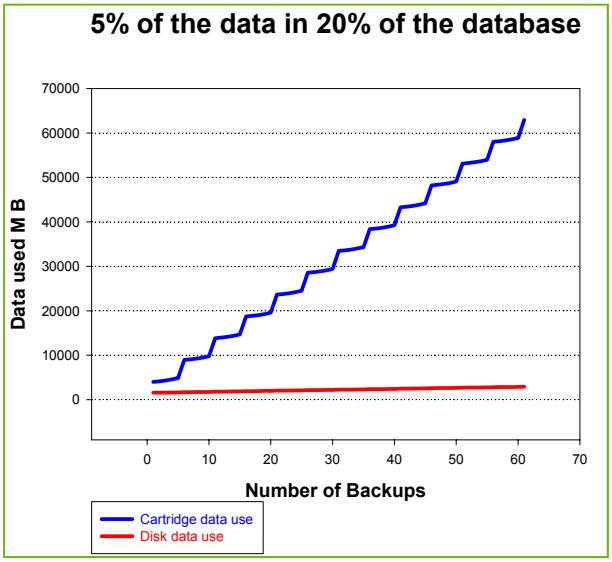

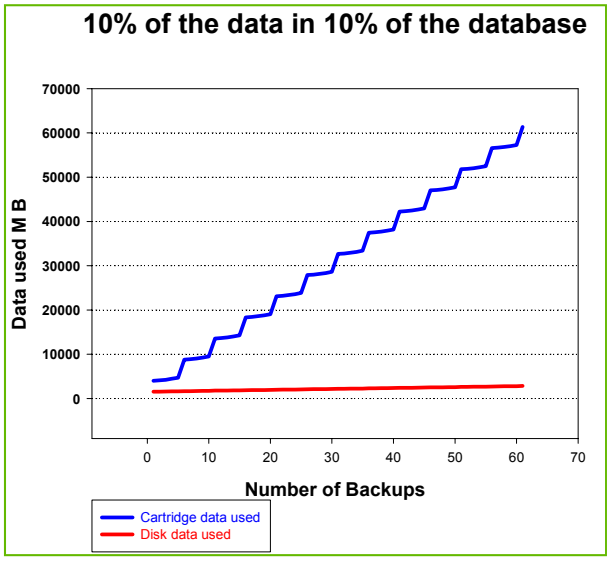

Actual savings in disk storage capacity after 3 months of testing

Test	Disk Data Usage MB	Cartridge Data Usage MB	Savings in MB
0.4% in 100%	2284	55374	53090
1% in 60%	2757	62694	59937
2% in 40%	2868	63261	60393
5% in 20%	2950	62942	59992
10% in 10%	2858	61344	58485

NB: Results for disk data used (red) appear as virtually flat lines on the graphs. This is due to the scaling required to incorporate the much greater values for cartridge data usage in the same graph. The figures shown in the table above confirm this.

Deduplication Test Results

Deduplication Ratios Using File Server test scenario

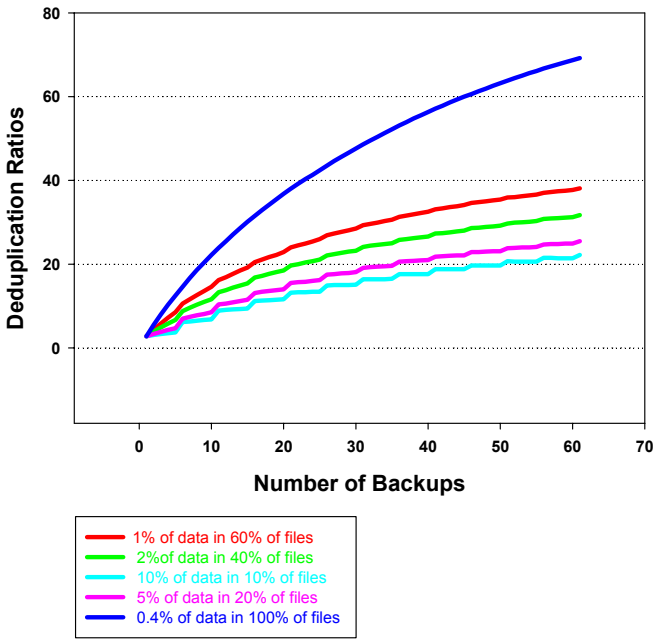

Deduplication Ratios Using Exchange Server test scenario

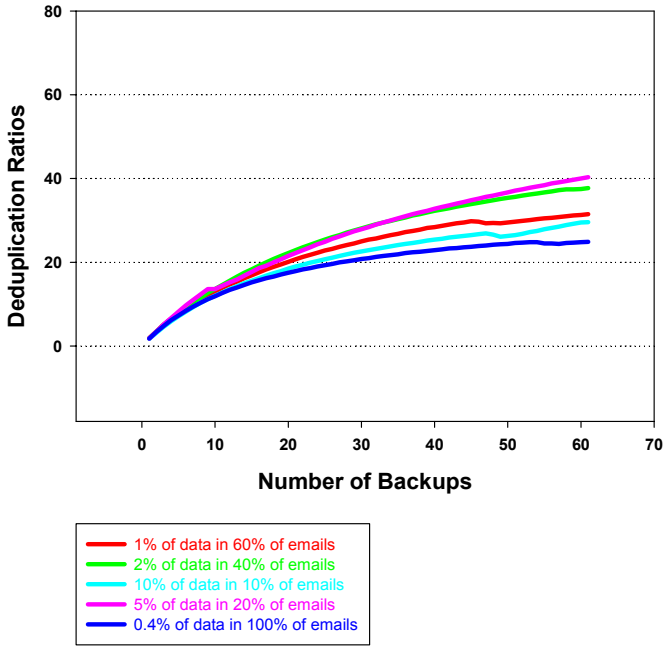

Deduplication Ratios Using SQL Server test scenario

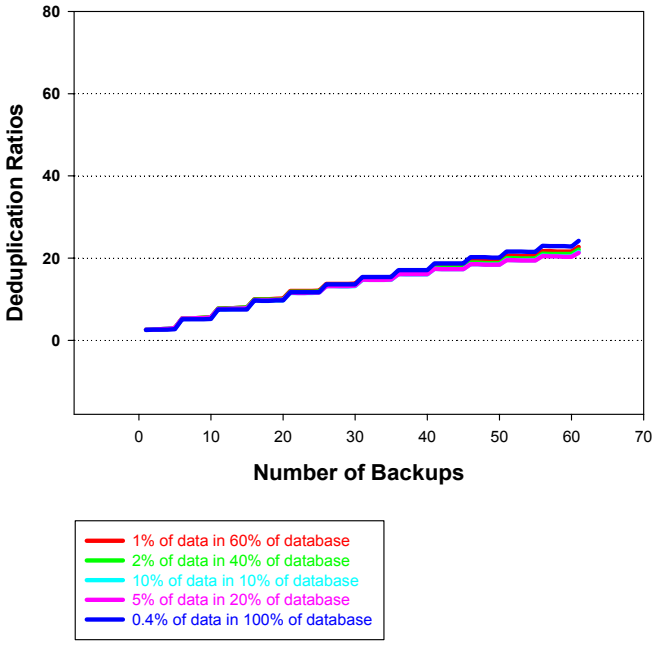

Maximum deduplication ratios after simulated 3 month backup period

Tests	File Server	SQL Server	Exchange Server
0.4% in 100%	69.2	24.2	24.9
1% in 60%	38.1	22.7	31.5
2% in 40%	31.7	22.1	37.7
5% in 20%	25.5	21.3	40.3
10% in 10%	22.2	21.5	29.6

Summary and Conclusion

The tests used in this report represent an experimental method which truly reflects real world scenarios and does not depend on extrapolations or projections based on limited testing. A key finding from our performance tests is that the deduplication technology behind the D2D4000 delivers consistent results across a range of applications. Furthermore, the results confirm that HP's claims of 20:1 to 50:1 deduplication ratios are quite achievable in the real world.

The best results were seen in the series of file server tests where greater opportunities for high deduplication rates present themselves. After the three-month test period had completed the best ratio we saw was nearly 70:1 although the 0.4% change in 100% of files test did effectively result in all daily incremental tasks securing the same amount of data as the full backups. However, the four other file server tests delivered deduplication ratios of between 22.2:1 and 38.1:1 which are within HP's claims.

The results from the SQL Server and Exchange Server tests also show a consistent rise in deduplication ratios over the three-month test period. However, the changes made to the test data occur at a more granular level so the end results are not as high as those seen in the file server tests. Nevertheless, for SQL Server the ratios range from 21.3:1 to 24.2:1 which are also within HP's claims. The twelve regular steps in the SQL Server Cartridge Data Used graph lines show when the full database backups are occurring and yet comparing these with the corresponding Disk Data Used graph line clearly reveal deduplication coming into play.

We were impressed with the high level of storage savings that were achieved with our 4GB test sample of data. At the end of the simulated three-month period the file server test realised capacity savings of between 64.9GB and 229.5GB. For Exchange Server this ranged between 156GB and 176.4GB whilst SQL Server returned savings of between 51.8GB and 59GB. Scale these results up to a full production environment with much greater backup needs and the potential for reducing storage requirements for backup is highly significant.

The D2D4000 Backup System delivers on HP's performance promises and we also found it very easy to install and deploy in our lab environment. The quick start wizard offers a fully automated setup routine that makes light work of initial installation allowing us to swiftly deploy the appliance in our test backup strategy. The main administrative web interface is also very intuitive and provides easy access to the various features. VTL creation is handled well and the appliance offers a good range of library and tape drive emulations along with support for either IP or FC SANs making it very versatile.

Faced with the inexorable growth in demand, administrators must look beyond traditional methods of data storage if they are to find a solution. Deduplication looks to have the answer and this report shows that measurable gains can be achieved by implementing it at the backup stage. It has the potential to dramatically reduce the amount of physical storage required for backup operations thus allowing more data to be kept within the data centre where it can be more easily accessed for restoration purposes.

Traditionally, deduplication has been a technology that only enterprises could afford but the D2D4000 Backup System bucks this trend and delivers it to the cost-sensitive SME. This appliance offers clear cost benefits in terms of reduced storage requirements, it is particularly easy to install and deploy and can have a big impact on management overheads by streamlining data protection strategies.

Copyright © 2008 Binary Testing Limited. All rights reserved. No part of this publication may be reproduced in any form without permission in writing from Binary Testing Limited.
Binary Testing Limited believes the information in this publication to be accurate as of its publication date and retains the right to make changes to this publication at any time without notification.
Binary Testing Limited makes no warranty with respect to the information that appears in this publication, assumes no responsibility for any errors that may appear and does not make a commitment to update the information contained therein.

All trademarks referenced in this publication belong to their respective companies

HP Publication number 4AA2-0799ENW

July 2008