HP StorageWorks D2D4000 Backup System

A report and full performance test of Low Bandwidth Replication on Hewlett-Packard's SME data deduplication appliances

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 deduplication is revolutionising the way businesses manage their storage resources and backup processes but many are now looking beyond these benefits to streamlining their disaster recovery procedures as well.

Compliancy with data protection regulations demands that secure off-site storage is maintained and replicating deduplicated disk based storage systems to remote locations is one of most efficient ways to achieve this.

Consolidating data deduplication and offsite replication into a single product has clear advantages and HP's latest D2D appliances deliver its Low Bandwidth Replication (LBR) technology which aims to offer SMEs major cost benefits over competing solutions.

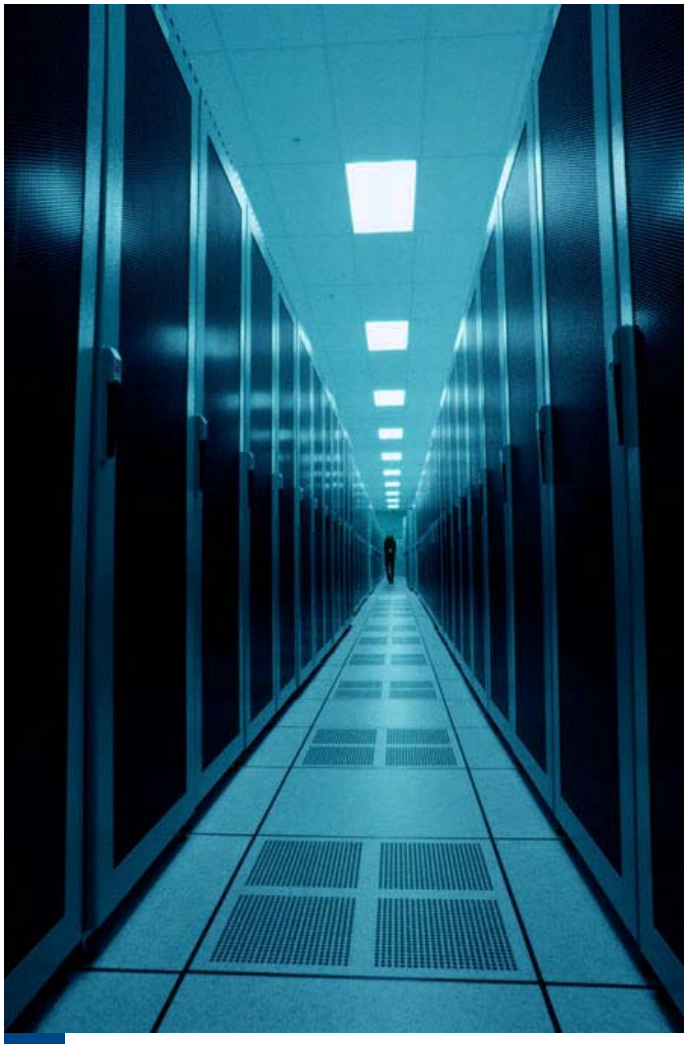

Businesses must have laid down disaster recovery procedures which include secure, off-site storage but many methods for achieving this are inefficient in terms of cost and management. The use of disk arrays has been growing exponentially as businesses see the benefits over solutions such as tape as a primary backup target. However, replicating these arrays to a secondary site without any form of data reduction inevitably requires expensive high bandwidth WAN links if current backup windows are to be met.

Teaming up its dynamic data deduplication and LBR technologies has allowed HP to address many of these issues. In tests run by Binary Testing in 2008, HP's D2D4000 appliance was shown to achieve measurable gains when implemented at the primary backup stage and the report confirmed that HP's deduplication dramatically reduces the amount of physical storage required for backup operations.

LBR takes advantage of these benefits as reductions in the capacity of primary backup storage allows it to replicate local virtual tape libraries (VTLs) to off-site locations within significantly smaller time windows. This has many benefits as businesses can designate fixed windows for replication and use low bandwidth WAN links to realise immediate operational costs savings.

Businesses with multiple remote offices replicating to a single centralised D2D appliance can use this reduced window to ensure the replication process only occurs outside normal working hours thus avoiding bandwidth contention with services during business hours. LBR can also support concurrent operations allowing many-to-one replications from remote sites or branch offices to occur simultaneously.

LBR fits well with businesses running 24/7 operations as once the local backup process has completed, replication can run in the background during the day but have throttling applied so it only uses a percentage of available bandwidth. It can reduce management overheads as the entire replication process is automated by the D2D appliances.

All replication is configured from the same web based GUI as used for deduplication so policies can be easily created and managed whilst wizards for replication and recovery reduce this to a simple four-step process.

HP also addresses the demands of full disaster recovery requirements that call for data to be archived from virtual tape to physically removable media. Its Tape Offload function uniquely allows data on virtual cartridges to be migrated to tape cartridges for long term storage and also for use in 'seeding' the remote appliances.

Executive Summary

This report provides a full evaluation and performance analysis of the D2D4000 backup appliances and focuses on validating HP's new low bandwidth replication (LBR) technology to determine the benefits it can offer for reducing the time required to replicate data off-site.

A range of tests has been used in order to determine how well LBR performs when automating off-site replication, the time savings that can be made for backup operations and what efficiencies can be introduced into WAN links.

As highly popular business applications run by the majority of SMEs, Microsoft's File Server, SQL Server and Exchange Server were chosen to run the tests on. For each application, the tests introduce controlled and measurable rates of change to a 50GB test data sample and run automated backups to standard practices covering a simulated period of one month. Two D2D4000 appliances were used with one storing local backup data from the test servers and the second acting as a remote site store for the automated replication operations.

The report provides a graphical analysis of LBR performance when replicating to an off-site appliance, the amount of time saved during replication operations and how efficiently the three test WAN speeds of 2Mbps. 10Mbps and 45Mbps were used.

The test results show clearly that considerable gains in a number of key areas can be made when using HP's LBR technology to offsite local backup data.

- LBR can make highly significant reductions in the time required to replicate data to a remote site.
- The reduced times will allow business to operate offsite replication in much smaller windows. This means they can be run overnight outside of normal working hours so negating any impact on general day-to-day operations.
- LBR makes more efficient use of the WAN link to the remote site allowing businesses to consider smaller connections for data offsiting operations which can result in large cost savings.
- The simplified installation and deployment processes for data replication make the D2D appliances a good choice for SMEs with limited IT support facilities.
- Combining the reduced WAN requirements of LBR with its blackout window and bandwidth limiting features will allow 24/7 businesses to run offsite replication tasks during normal working hours with minimal impact on general IT operations.
- As all LBR operations are fully automated they are far more reliable than the common practice of manually shipping tapes offsite. Further time savings in man hours can also be made as support staff are freed up from these tasks.

Average time (hours) to replicate 50GB over a 2Mbps WAN link

	Seeding runs [1]		Full Backups		Mid-week backups	
	With LBR	Without LBR [2]	With LBR	Without LBR [2]	With LBR	Without LBR [2]
SQL 2% in 100% of rows	11.1	54.8	0.6	54.8	0.6	4.5
SQL 5% in 100% of rows	11.1	54.8	0.9	54.8	1.5	11.0
Exchange 5% in 100% of emails	46.4	62.5	1.5	62.8	0.2	2.6
File Server 5% in 10% of files	3.1	57.3	0.5	57.3	0.1	5.7
File Server 10% in 10% of files	3.1	57.3	0.5	57.3	0.3	5.7

[1] The seeding runs refer to the initial transfer of data between the local and remote D2D appliances.

[2] Without LBR refers to the calculated minimum time taken for replication to complete without any form of data reduction or compression

The HP StorageWorks D2D4000 Backup System

HP is currently offering three solutions for the SME market with the family comprising the D2D2500, the D2D4000 featured in this report and the D2D4100 which offers up to 18TB of usable capacity and can support a 24-to-1 fan in from appliances located at remote sites.

The D2D4000 is a 2U rack mount appliance with room at the front for up to twelve drives with the base 4.5TB model starting with six 750GB SATA hard disks. The systems under test are both the 9TB models and were supplied to us with twelve 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.

You can choose between dual Gigabit Ethernet ports for iSCSI operations and dual 4Gbps fibre channel ports for integration into an FC SAN and the appliances support both of these simultaneously. Both appliances are shipped with the drives configured in a RAID-6 array which provides dual drive redundancy.

In either IP or FC SAN configurations the appliances can present up to sixteen VTL targets allowing them 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 virtual tape drives. Depending on the model chosen, the D2D4000 appliances can emulate up to 24 virtual tape devices and support up to 3,456 virtual tape cartridges

HP offers a simplified and cost-effective licensing scheme as Dynamic deduplication is included as standard. For LBR only a single license is required for the replication target and once activated it can be used immediately by as many sources as the target appliance model can support.

Removable storage requirements are covered as HP offers optional kits comprising LTO-2, LTO-3 or LTO-4 external tape drives and SCSI or SAS host bus adapter cards. Drives can also be direct attached to the D2D4000 appliances to allow the Tape Offload and Import functions to be used for the initial seeding runs.

HARDWARE SPECIFICATION

Chassis: 2U Rack Mount	CPU: 2 x 3GHz dual-core AMD Opteron 2222
Memory: 8GB 667MHz DDR2	Storage: 12 x 750GB SATA hard disks in RAID-6 array
RAID: HP Smart Array P400 SAS/SATA with 512MB cache and BBU	Network: 2 x Gigabit Ethernet
Power: 2 x 750W hot-plug supplies	Management: Web browser

binary testing

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.

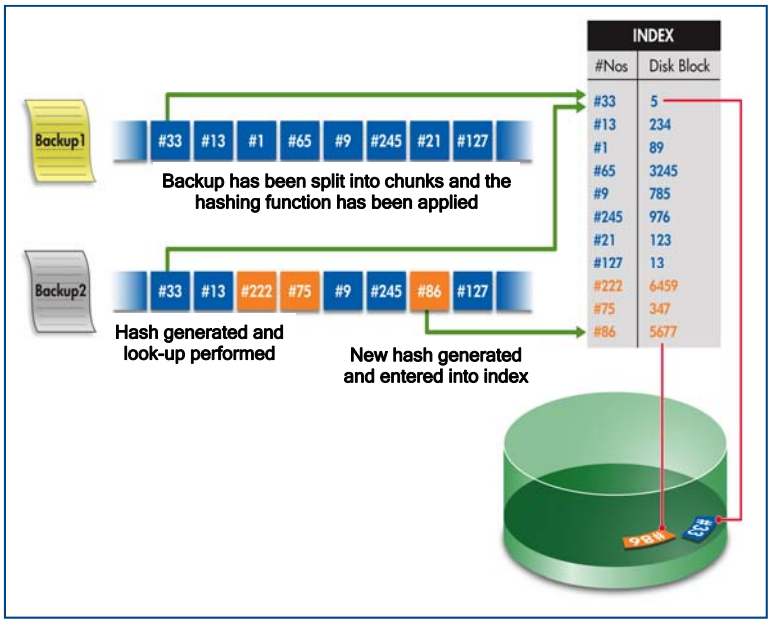

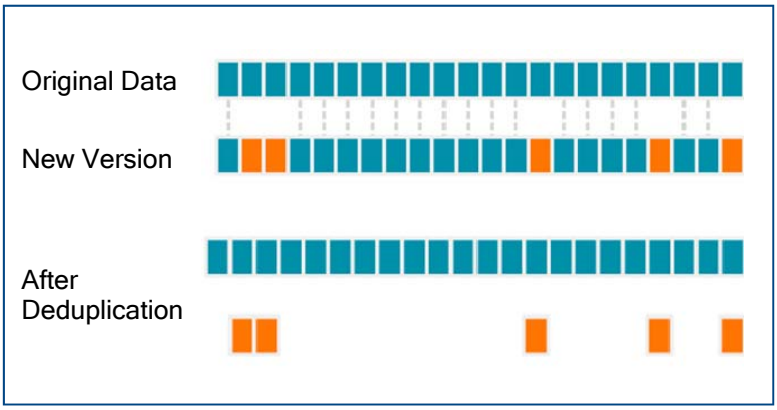

The D2D4000 uses a nominal chunk size of 4k to compute its hash values. To improve performance the index is stored in RAM but even on the D2D4000 with 3TB of usable storage the requirements could be as high as 30GB of memory which is clearly cost prohibitive.

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.

A transaction logging system ensures that there is no data loss in the event of a power failure. In response to virtual tape events such as 'write filemark' and 'end of data', the critical data stored in memory is flushed to disk.

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 backup data format.

HP's Low Bandwidth Replication

HP's LBR synchronises data between source and target appliances so that they are an exact match. It works at the cartridge level where you choose the slots in the source VTL you want replicated and these are mapped to corresponding slots in the target VTL at the remote site

After the first full backup is taken, the target VTL needs to be 'seeded' with the data from the source but on subsequent backups only new data needs to be replicated. This comprises any unique 4K data chunks along with their hash codes plus control, or 'manifest' data.

The D2D appliances support three seeding methods:

1) Seeding can be carried out directly between the local and remote appliances over the WAN link.
2) The source and target appliances can be initially co-located on the main LAN and seeding carried out over a high speed Gigabit connection. On completion, the target appliance is then shipped to its remote location
3) An LTO tape drive or library is connected to the local appliance and the virtual tapes copied to physical media. These are then shipped to the remote location where they are used to create new virtual tapes on the replication target appliance. On completion, the target appliance is then synchronised with the source appliance over the WAN link.

These seeding methods provide businesses with a wide choice of initial deployment methods. Methods 2 and 3 are most likely to be used but our LBR performance results show that in many cases it is quite feasible to carry this out directly over a low bandwidth WAN link.

Method 3 will be used where the remote appliance is already in situ and is achieved using a function called Tape Offload. This is a standard feature of the D2D appliances which allows data on virtual cartridges to be copied to physical tapes. These are then shipped to the remote site and from the appliance's browser management interface are imported to a new virtual cartridge.

D2D Replication Deployment Scenarios

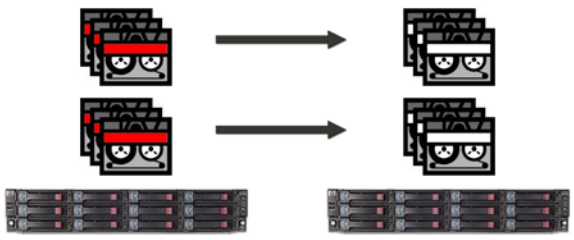

ACTIVE– PASSIVE CONFIGURATION

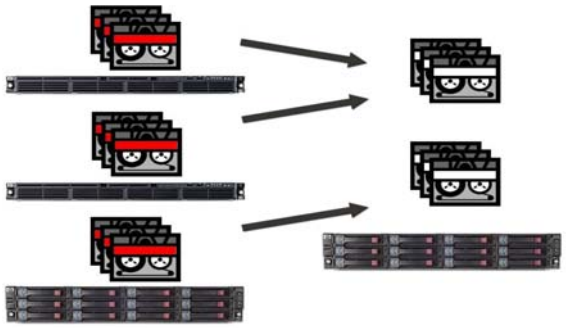

MANY–TO–ONE CONFIGURATION

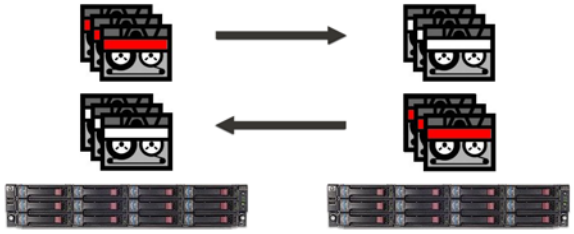

ACTIVE-ACTIVE CONFIGURATION

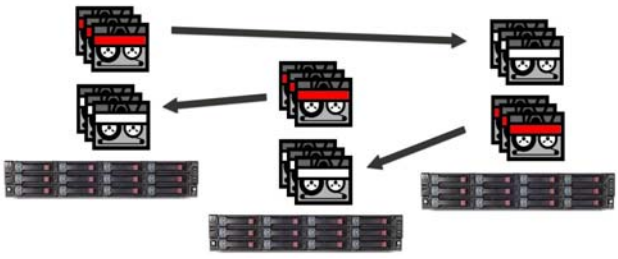

N-WAY CONFIGURATION

LBR supports a wide range of deployment scenarios making it extremely versatile. A simple setup would be an active/passive configuration where a number of source VTLs on a single appliance are replicated to a remote VTL but it's also possible to use an active/active setup where VTLs on both sides are being replicated.

Target VTLs can accept replication from up to four different sources and each target appliance supports multiple VTLs. Called Fan Out, this allows you, for example, to back up data to a number of source VTLs and replicate each one to different remote locations. The D2D appliances currently support a maximum of four targets for Fan Out operations.

Many-To-One, or Fan In, can be used to replicate multiple remote VTLs to a single, centralised source. This will be invaluable for businesses that have local backup operations being conducted at remote or branch offices and want to provide essential disaster recovery for all of them. The D2D4000 family supports up to sixteen source appliances which has been extended to twenty-four on the new D2D4100 appliances.

The data recovery options offered by the D2D appliances are extensive and where a file or folder needs to be reinstated from a VTL the methods are exactly the same as for a standard physical tape library. Furthermore, as virtual tapes load faster and can be searched much more quickly, data restoration performance is greatly enhanced.

In the event of a remote site disaster possibly involving total systems and data loss, there are a number of recovery options available.

1) Where a remote site is replicating to a target at the data centre or head office, by removing the relevant replication slot mappings, the target appliance can be converted into a non-replicating library. A new server is configured, connected to the appliance and data on the VTLs is restored to the server. The new server and D2D appliance can then be shipped to the remote site where LBR operations are then re-established.

2) The second method of data restoration is called 'reverse replication'. A new source appliance is shipped to the remote site and restored by replicating data to it from the original target appliance. Essentially, the new appliance is being seeded over the WAN link. On completion, the server at the remote office can be fully restored from the new appliance allowing it to perform normal backup operations again after which replication can begin again.

3) The third method uses the Tape Offload function where data from the target appliance at the head office is copied directly onto physical tapes. These are then shipped to the remote site along with the tape library or drive they were prepared on and imported into a new appliance. The data on the physical tapes is fully compatible with the backup application format and the import process regenerates all the necessary hash codes and the deduplication store.

Initial installation for deduplication operations is simplified by wizard-based routines for all configuration processes. For host systems, these install the LTO tape drivers, configure iSCSI services and discover D2D appliances on the network. During testing we found replication just as easy to configure as the entire process is conducted only at the source appliance and can be completed in a few minutes.

Usefully, when an iSCSI initiator connects to the D2D appliance for the first time a VTL is automatically created 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.

During testing we found replication very simple to configure as this is all conducted at the source appliance.

From the very intuitive web interface a wizard guides you through four key steps where you select a source VTL, provide the IP address of the remote appliance, view its VTLs and choose a target to replicate to.

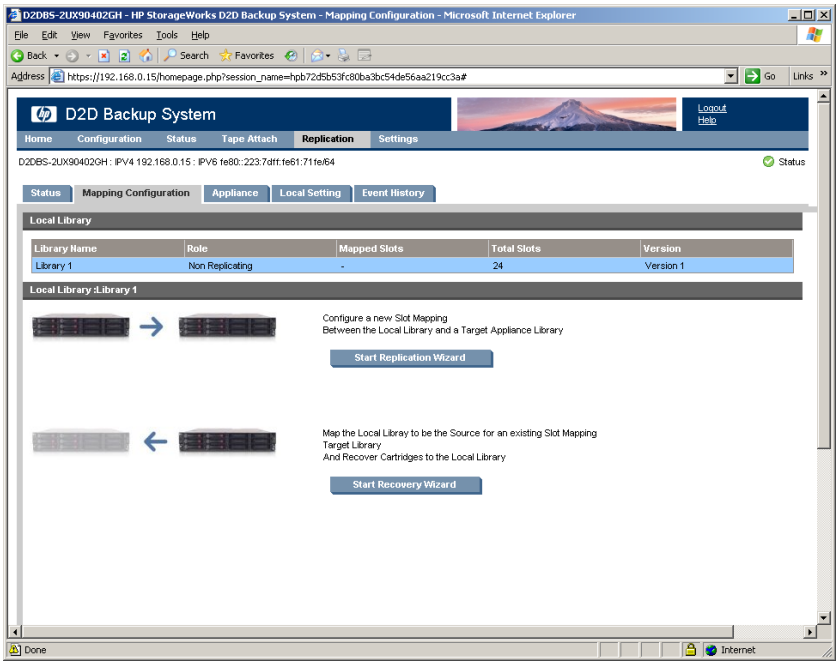

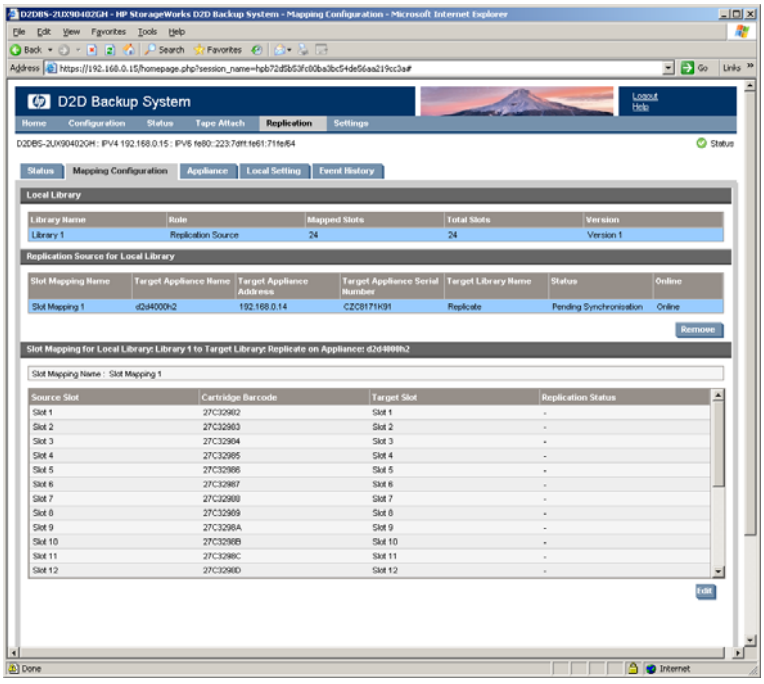

You then set up slot mappings and this defaults to mapping all slots in the source VTL to available slots on the target.

During this process you can use an existing library on the target for replication or create a new one from the source where all cartridge barcodes will be identical on both VTLs.

There is a fifth step that provides access to non-essential configuration. You can designate a black-out window during which replication is not permitted to run and select different start and finish times for each day of the week.

Bandwidth controls are also configured from here where you specify the link speed and decide how much WAN bandwidth the replication can use.

And that's all there is to do as dependent on any blackout window restrictions, replication tasks start as soon as the source appliance unloads a cartridge.

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 plus the SQL Server agent and connected to the local D2D4000 appliance over iSCSI. 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 72GB 15K SAS in RAID-1 mirror

To test HP's LBR technology we created a performance testing suite that mirrored real world scenarios closely. We devised a group of tests that allowed us to manipulate changes to the data held on the test servers at a known rate and quantity prior to it being backed up and replicated. These changes were applied to a Windows File Server, a Microsoft SQL Server database and an Exchange Server. The tests simulated a one month backup period where the data from all backups for each test sequence was retained on the local appliance.

For the File Server test we called for X% of data to be changed in Y% of files, for SQL Server test this was X% of data changed in Y% of database rows whilst for Exchange Server this would be applied to attachments sent with individual emails. Two combinations of data rate changes were applied to File Server and SQL Server and one rate change made to the Exchange Server.

The SQL Server database held 50GB of data in tables whereas the Exchange Server had 1,000 users with message attachments totalling 50GB distributed amongst them.

Data rate changes		
Data Type	**X% of data changed**	**In Y% of files, emails or database rows**
Exchange Server	5	100
SQL Server	2	100
	5	100
File Server	5	10
	10	10

WAN Speeds		
2Mbps	10Mbps	45Mbps

For the replication part of the performance testing we used one D2D4000 as a local backup appliance and a second to act as a remote site appliance which would function as the target. In between the two we placed a Network Nightmare WAN simulator appropriately configured for the required link speed. The tests were run over three WAN link speeds – 2Mbps, 10Mbps and 45Mbps.

The backup strategies for both applications 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

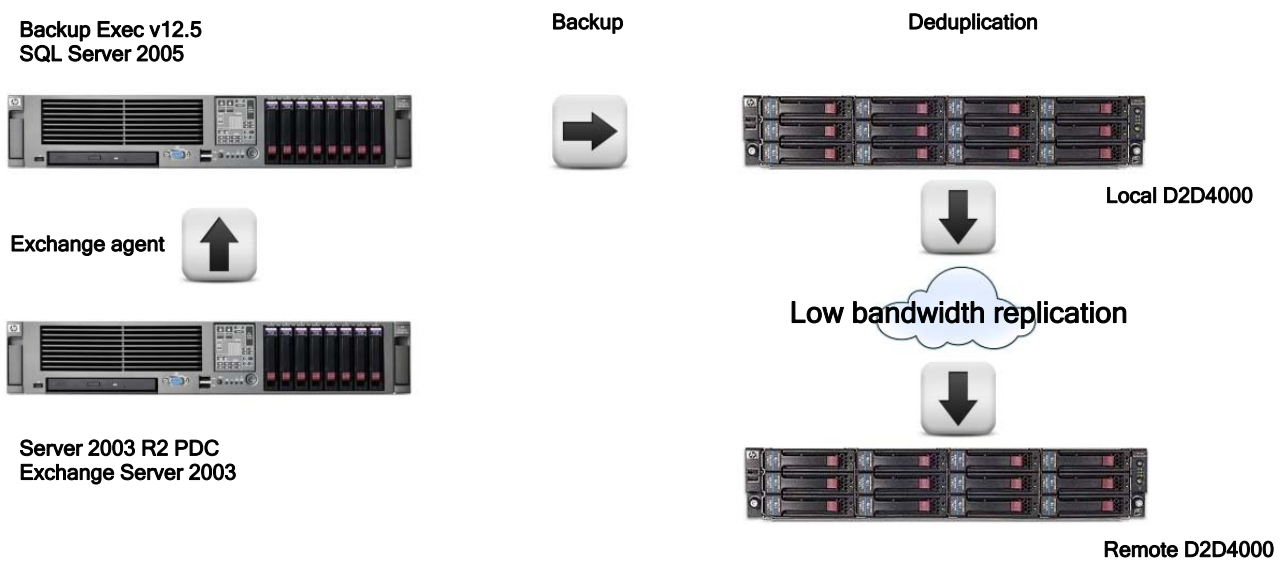

Our automated testing suite was designed to manage all aspects of the test process, namely amending the data sets, controlling the type and execution of the backup processes and then extracting the replication performance data from the D2D4000 devices. The suite was written using a combination of Visual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

To ensure the tests mirrored real world applications we collected data from archives of real systems and used this to create the test data sets. 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 wrap around to the beginning and continue or the simulated one month test cycle had concluded.

From the local appliance, we configured replication to start as soon as each backup job had finished and the virtual cartridge had been unloaded. Replication was monitored continuously with automated scripts and backup jobs were only allowed to start on successful completion of each replication task. No time restrictions or bandwidth limitations were specified allowing replication to have full use of the simulated WAN link.

It was also important that each tests' 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.

Using our automation scripts, storage usage data on the local appliance was collected after every backup had completed and inspected to ensure the deduplication processes were functioning smoothly. The scripts recorded the times for each replication tasks allowing us to calculate information on the actual and effective WAN speeds that were being delivered by LBR.

For SQL Server, the mid-week backups were effectively running two separate jobs – a transaction log backup followed by a differential. This resulted in two replication tasks being run and for reporting purposes we recorded an average rate for both jobs.

Replication times—Initial seeding runs

The following graphs show the times recorded for the initial seeding runs to the remote VTL as compared with the times they would have taken without HP's LBR. The latter times were calculated using the amount of data for the initial full backup as reported by Backup Exec for each job and dividing it by the absolute speed of the test WAN link.

Inefficiencies such as protocol overheads and bandwidth contention were not factored in as there are no definitive measurements for these. Consequently, it can be safely assumed that the performance improvements of LBR will be even greater in the real world.

As can be seen, significant time savings can be made with the initial seeding runs as data deduplication and compression are dynamically applied during the first full backup to the local appliance.

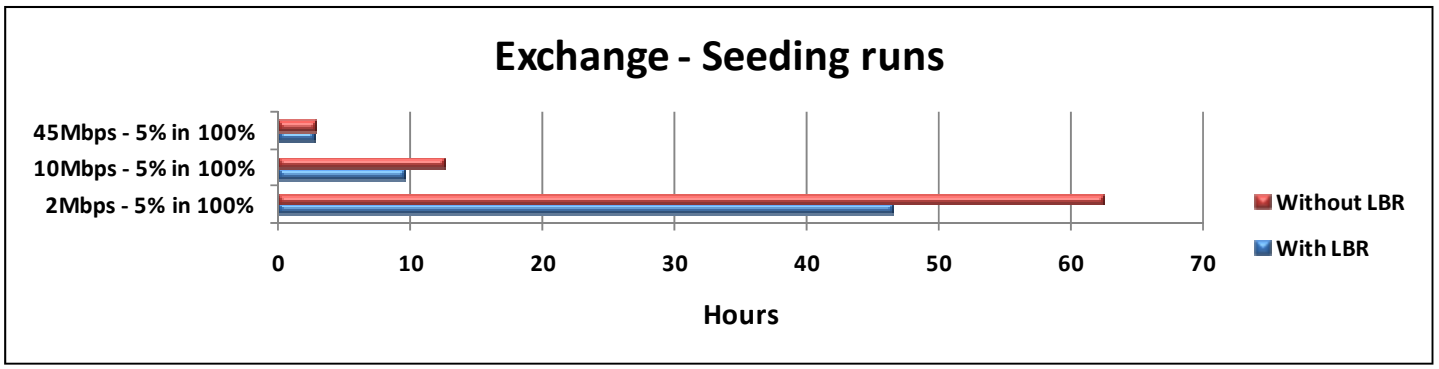

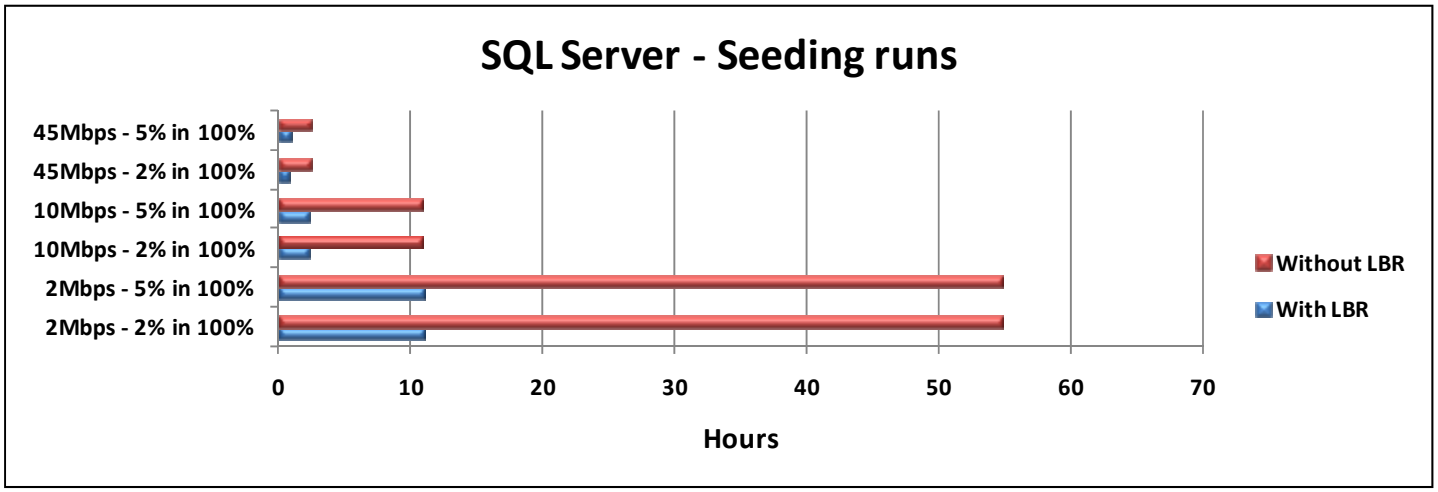

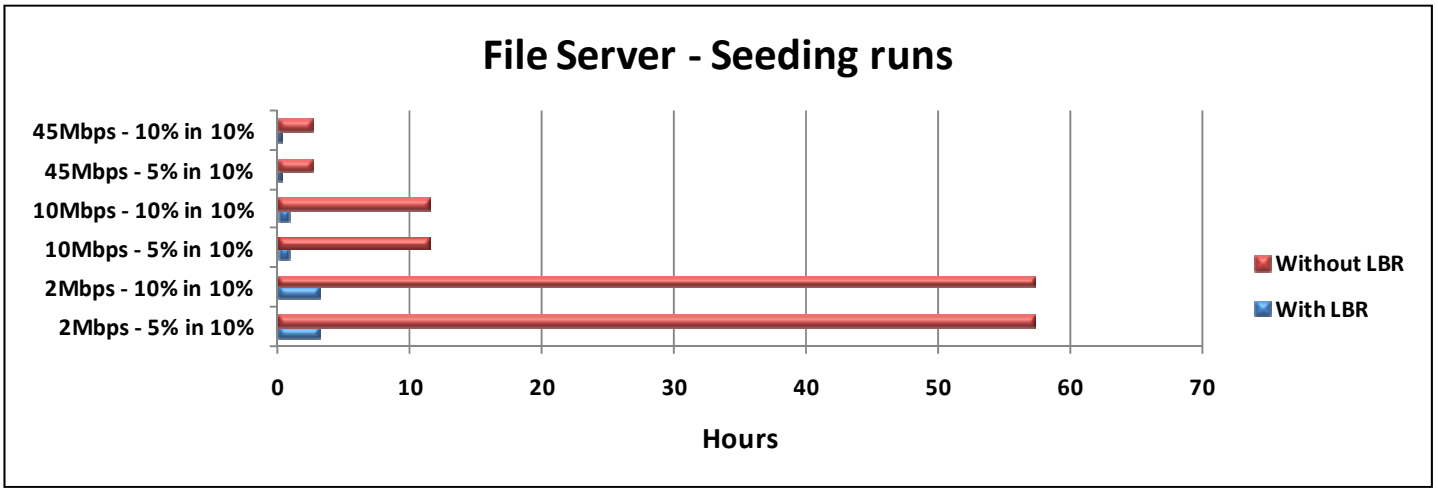

Replication times—Full backups

The following sets of graphs show the average times taken to replicate the four end-of-week full backups over the one month testing period. These are compared with the times replication would have taken without HP's LBR

The initial seeding runs have been excluded (see graphs on p.11) and the calculations for the theoretical times use the absolute speed of the WAN link and have not factored in any link inefficiencies.

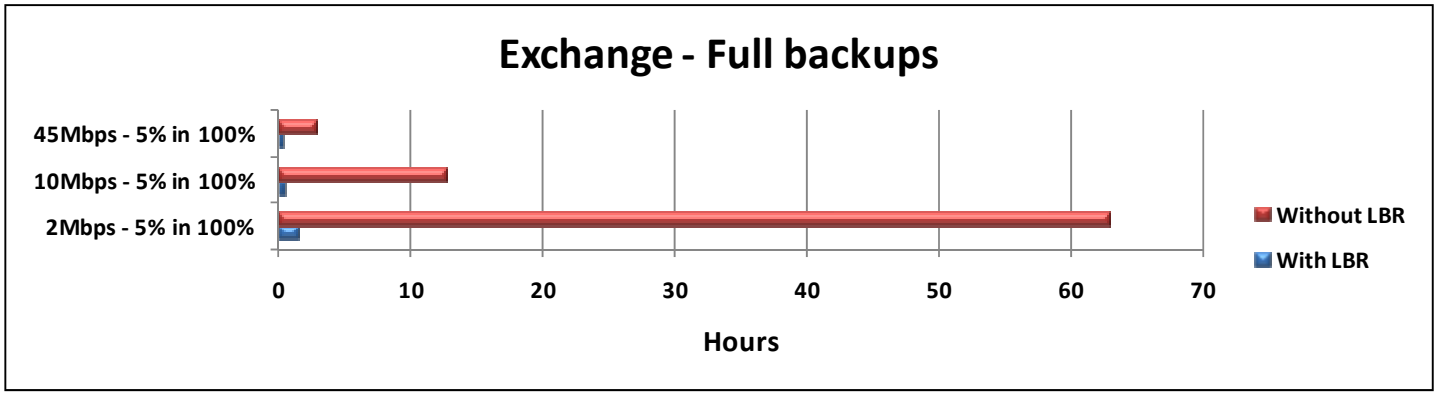

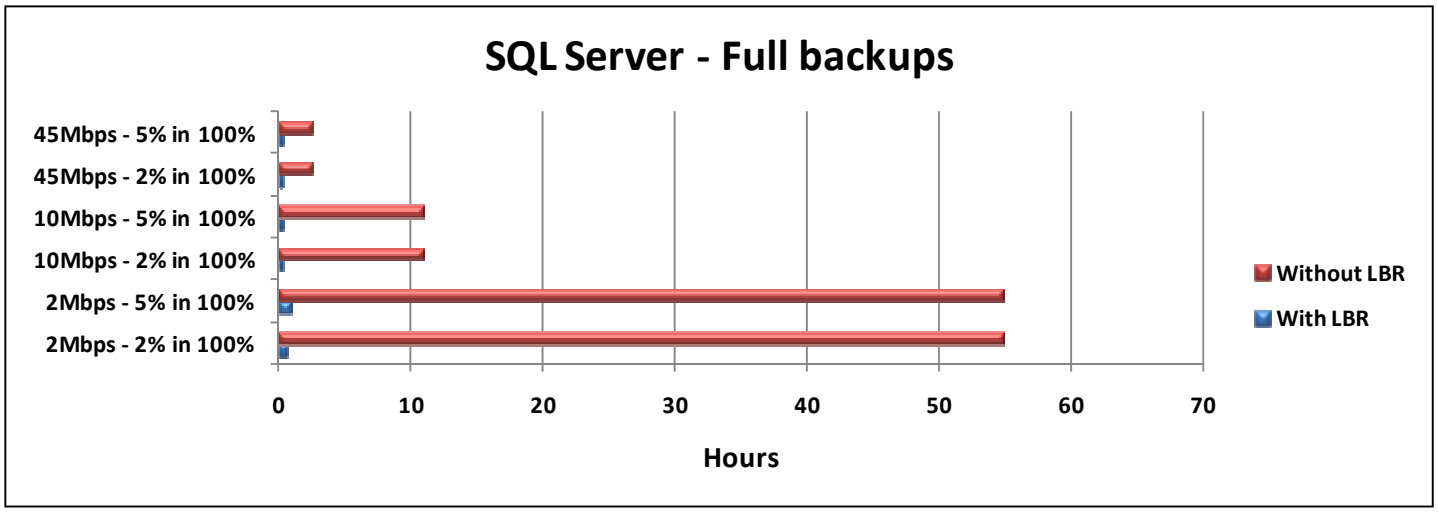

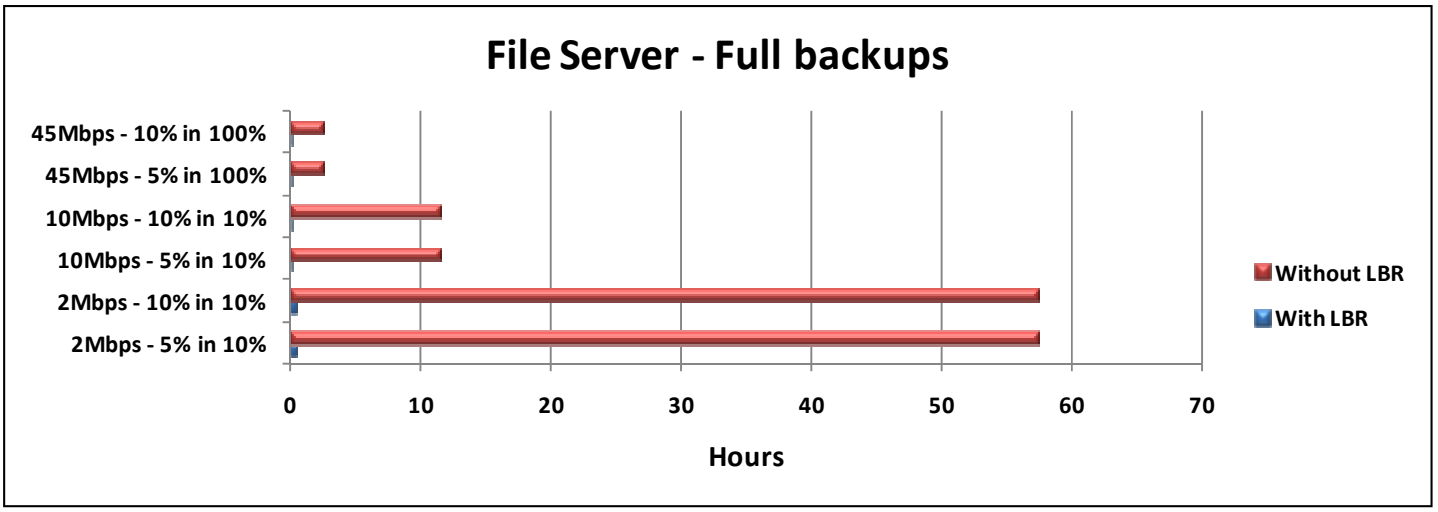

Replication times—Mid-week backups

The following sets of graphs show the average times taken to replicate the mid-week backups over the one month testing period. These are compared with the times replication would have taken without HP's LBR.

Note that the times taken with LBR for some of these runs were so short that the values barely show on the graphs when compared to the times without LBR.

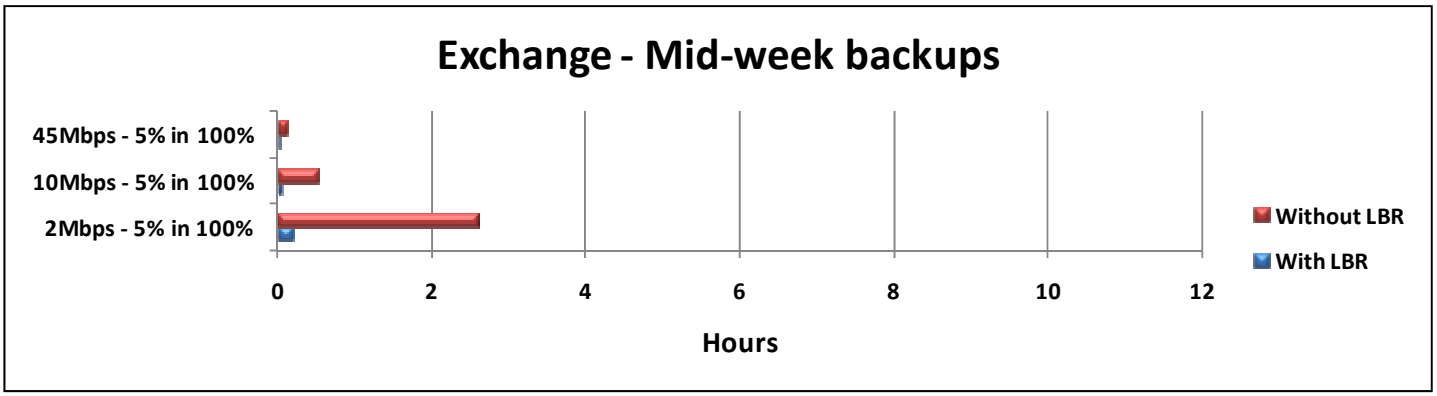

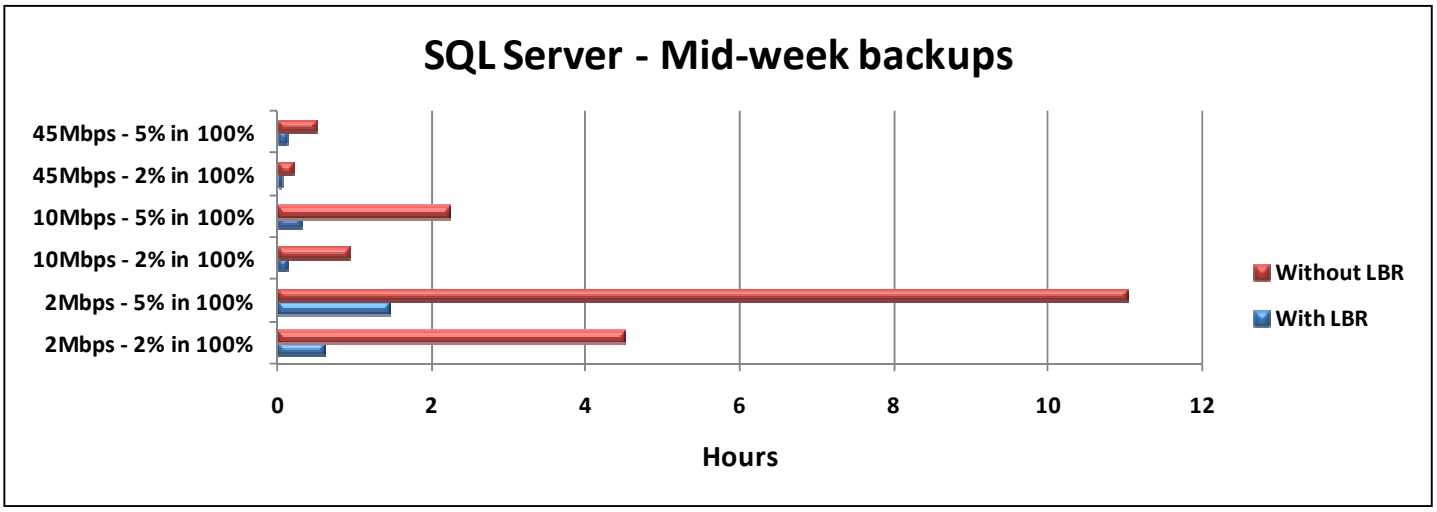

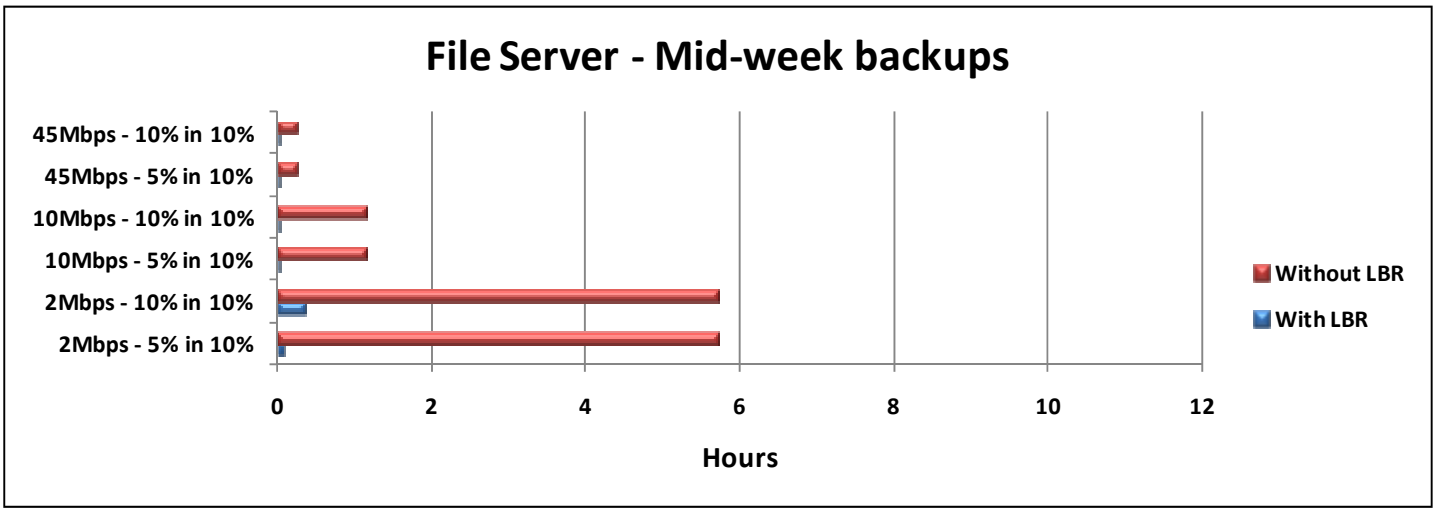

Backup Time Savings with LBR—2Mbps WAN

The following graphs show the overall time savings for replication operations that were made with LBR. The smaller area in blue shows the cumulative times for replication with LBR over the one-month period.

The red areas show the cumulative time the replication tasks would have taken without HP's LBR. These were calculated using the amount of data being backed up as reported by Symantec Backup Exec for each job and dividing it by the absolute speed of the configured WAN link.

The initial seeding runs were not included in these results. It should also be noted that the actual times taken for some runs with LBR were negligible when compared to the times without LBR and so hardly show on the graphs.

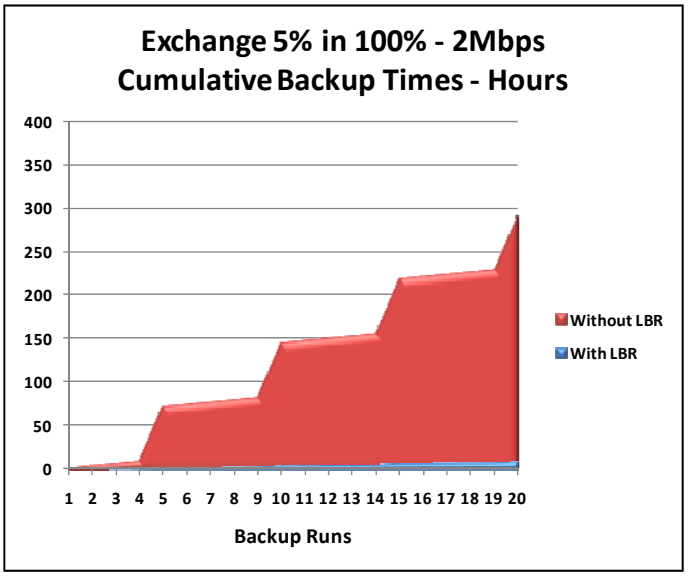

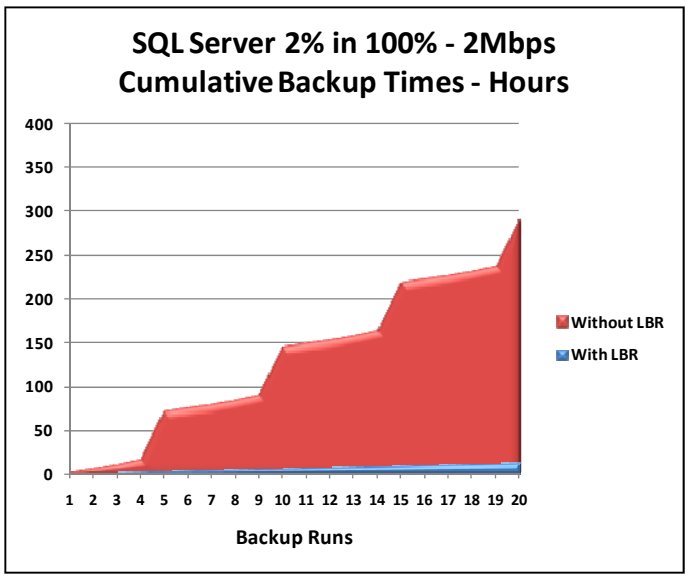

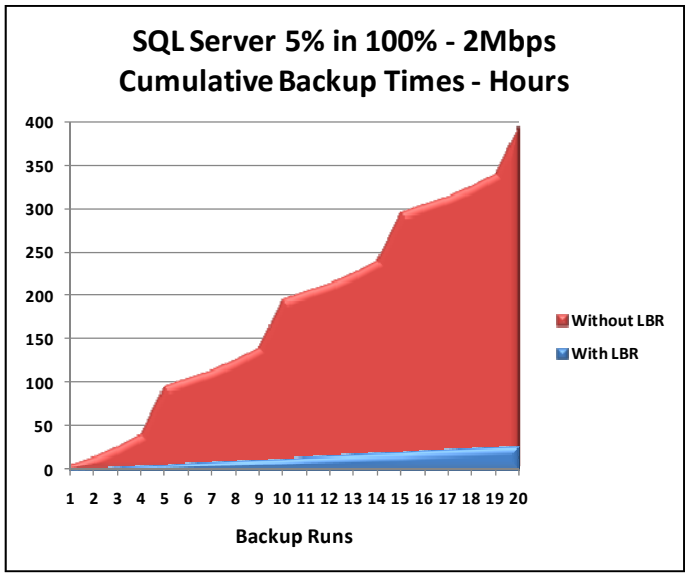

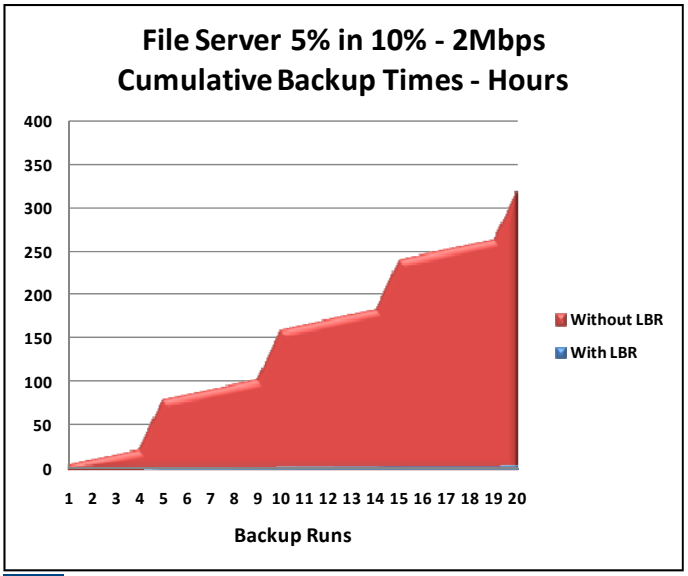

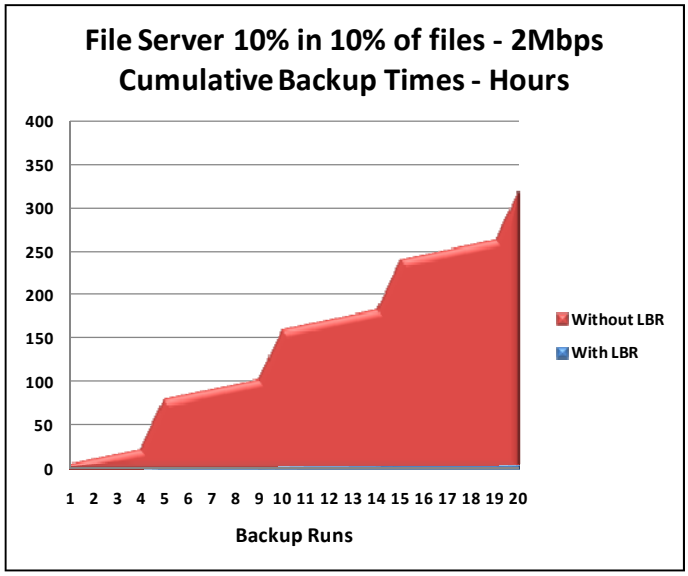

Backup Time Savings with LBR—10Mbps WAN

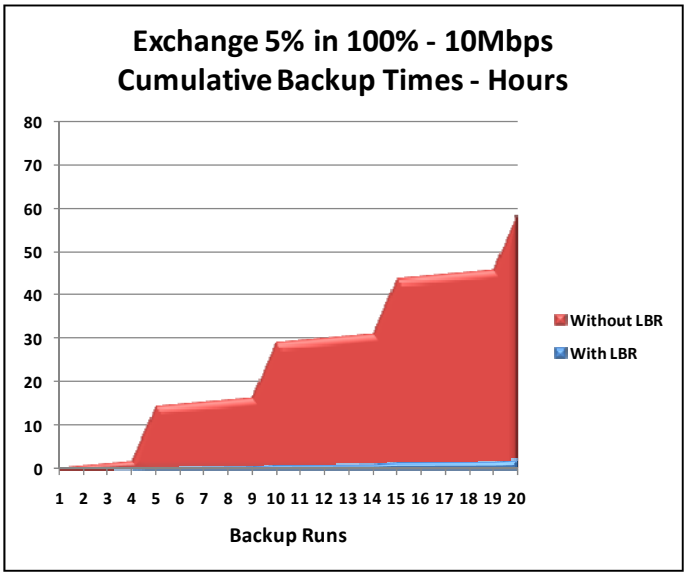

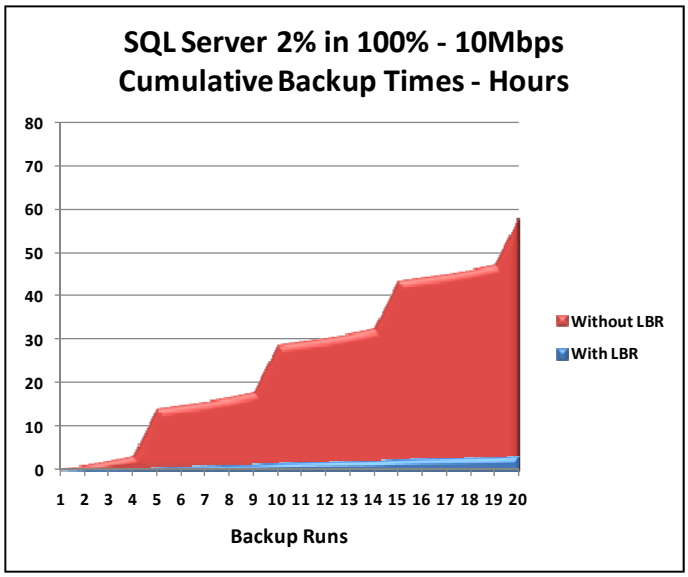

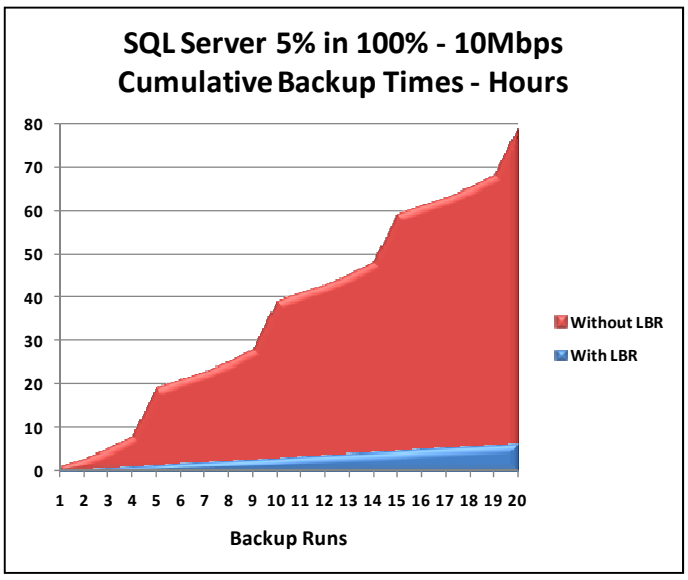

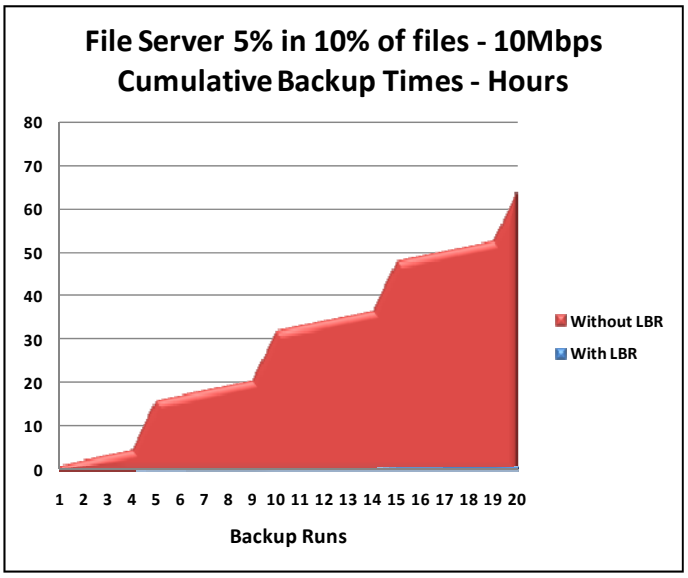

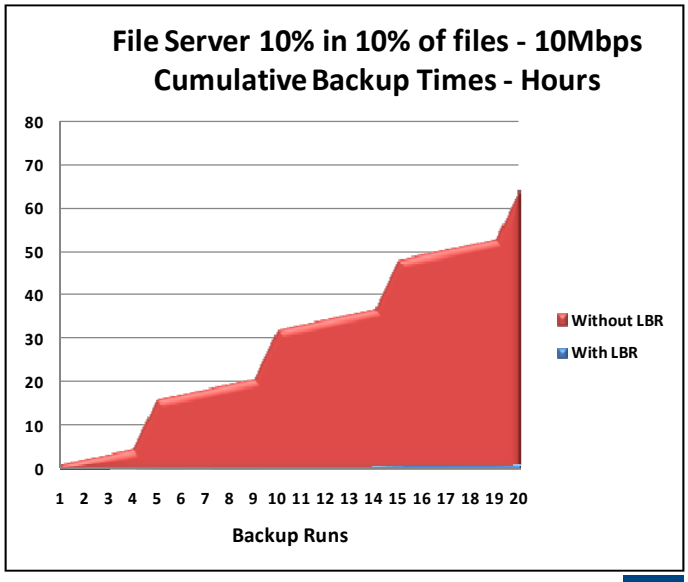

Backup Time Savings with LBR—45Mbps WAN

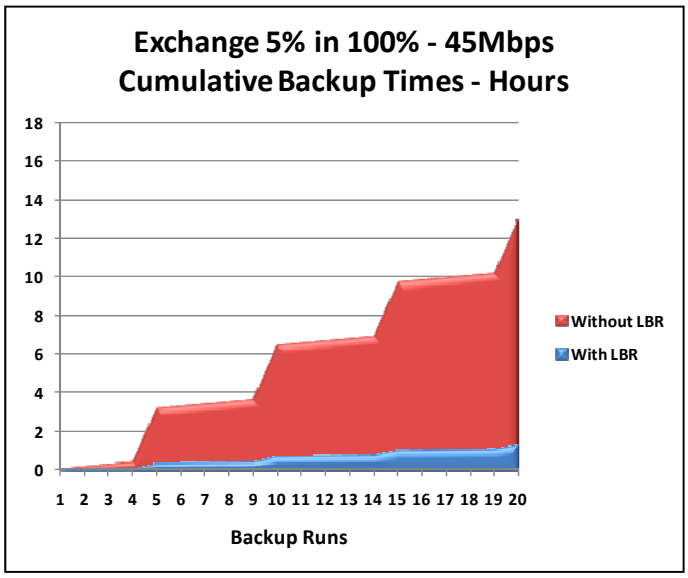

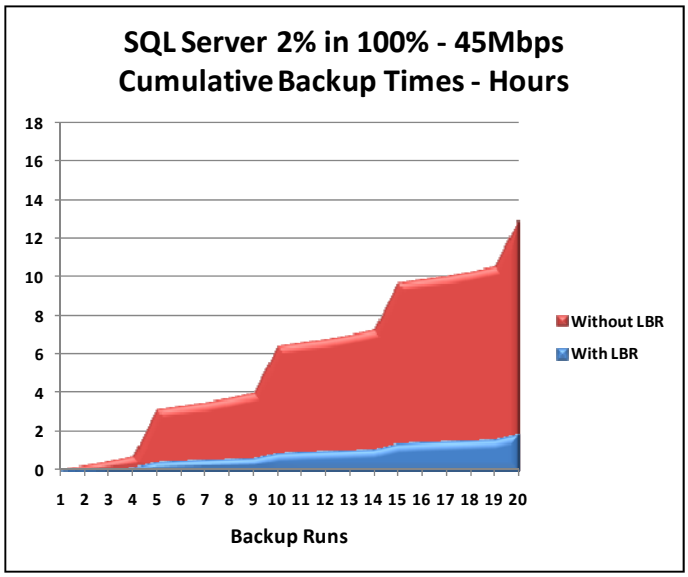

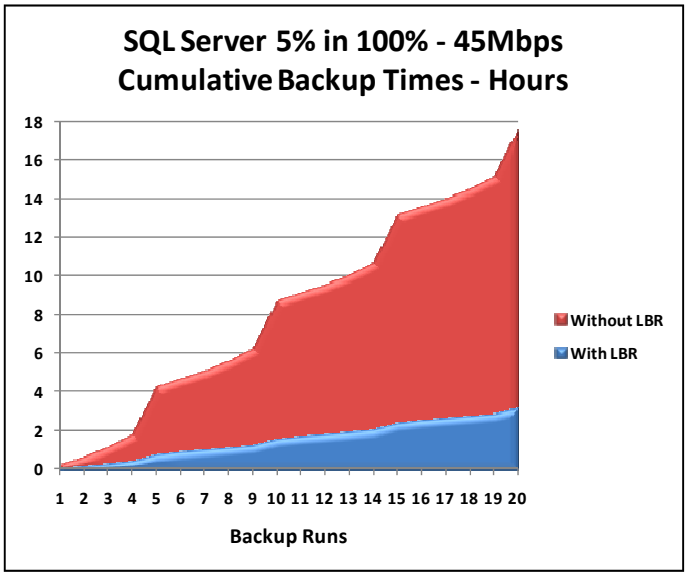

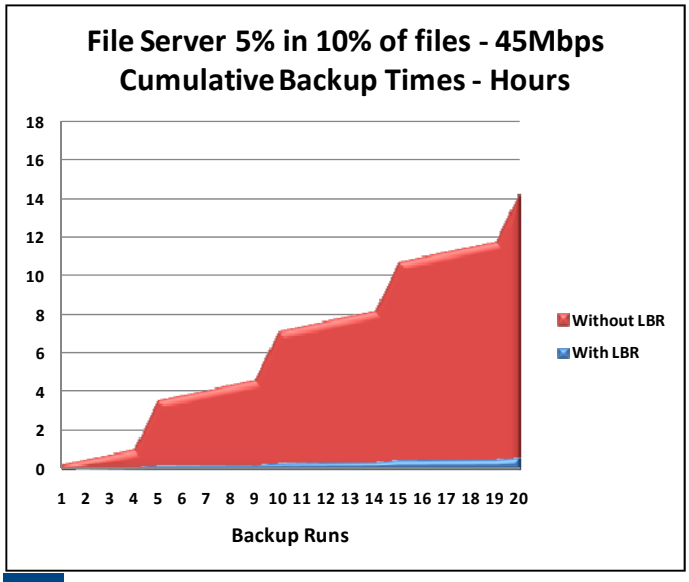

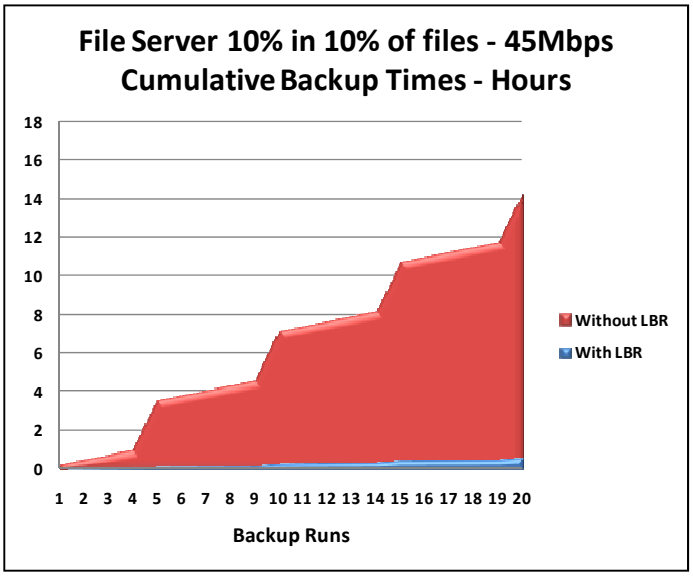

Apparent WAN Speeds

The apparent WAN speeds are derived from the amount of data being backed up as reported by Backup Exec for each job and dividing it by the actual time taken for each replication task. The results shown in the graphs below are an average apparent speed of all tasks for each data set and exclude the initial seeding run.

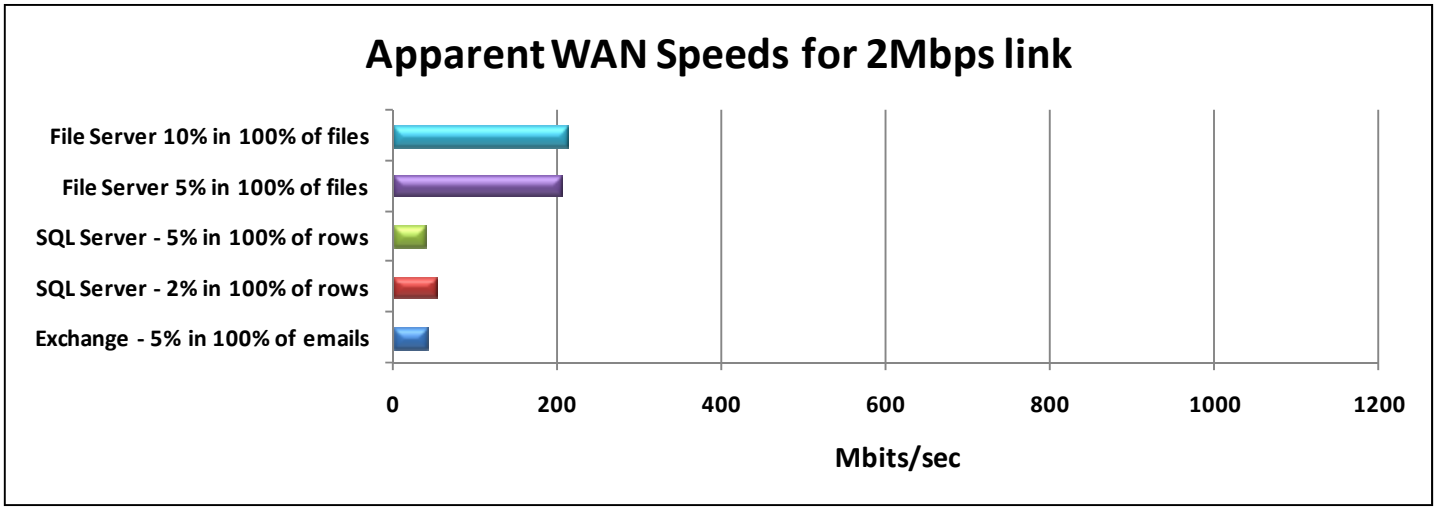

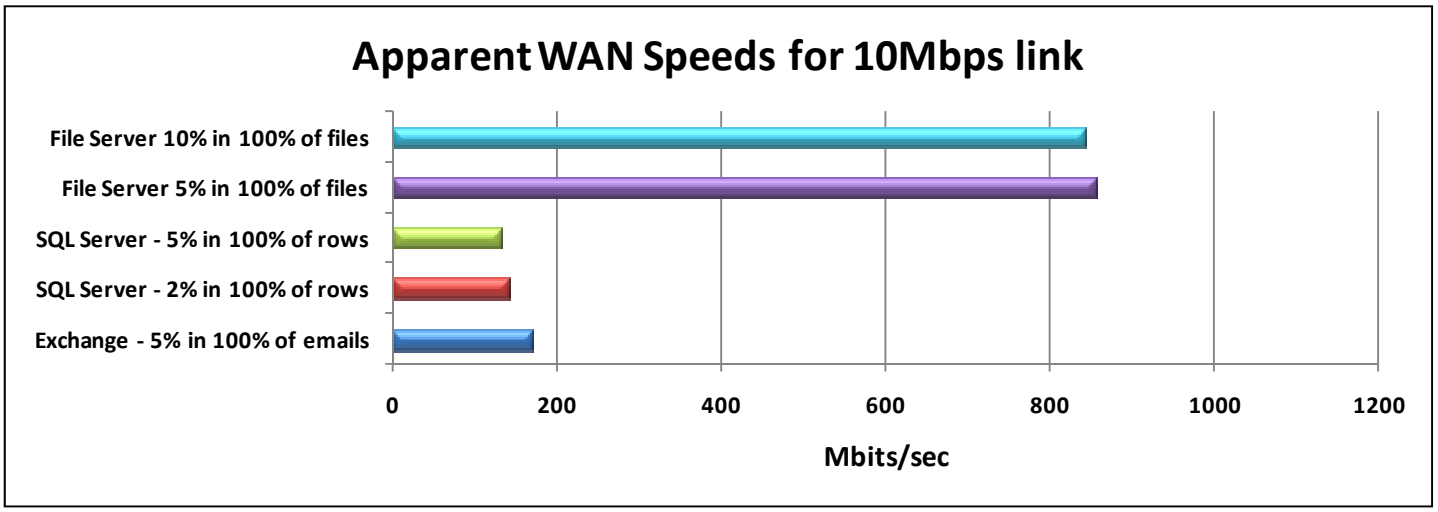

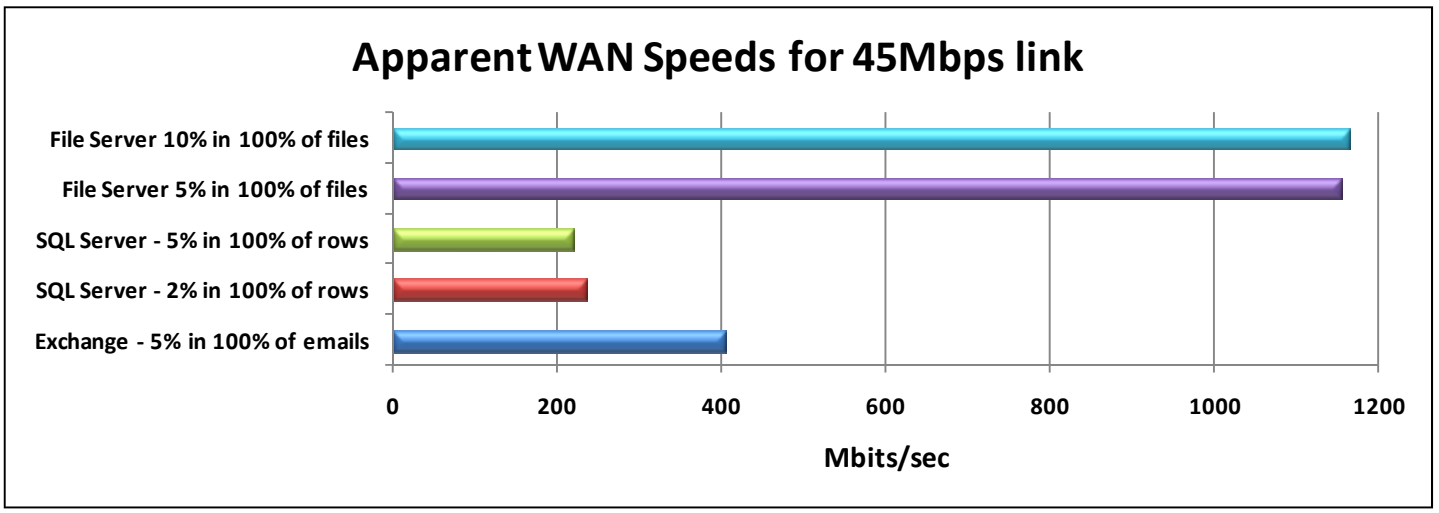

Many-to-One Replication Test

Many-To-One Replication Lab Test Environment

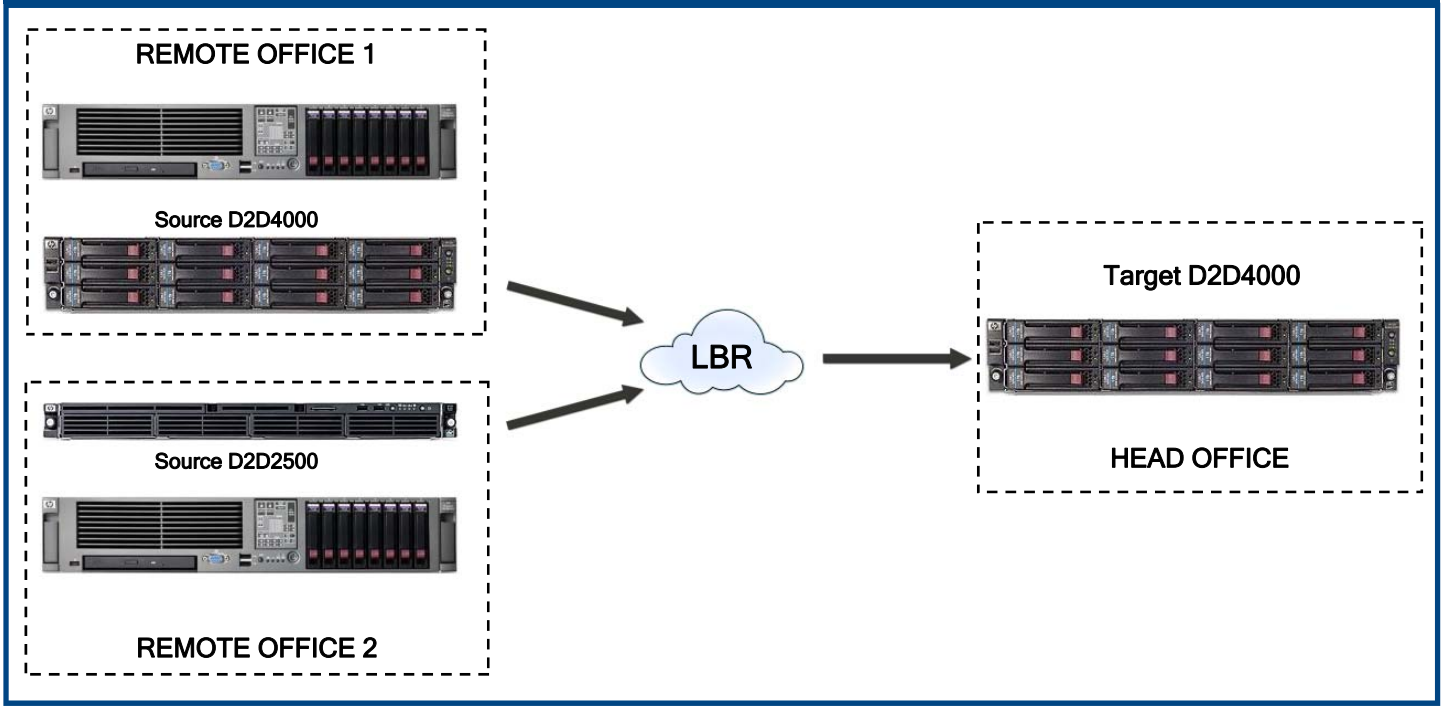

To test Many-To-One, or Fan-In, replication operations, we created an environment with two remote offices, each with their own Windows server carrying out local backups to dedicated D2D appliances.

Remote Office 1 had a Windows Server 2008 file server using Symantec Backup Exec 12.5 to secure the 50GB test data sample to a local D2D4000 appliance. Remote Office 2 had a Windows Server 2003 R2 system also providing file services and using Backup Exec to secure the same data sample to a local D2D2500 appliance.

The two D2D appliances were configured to replicate to a single D2D4000 appliance acting as the head office storage location for the two remote offices and connected via a 45Mbps simulated WAN link. An initial full backup of the 50GB file server test sample was run at each remote office to the relevant D2D appliance.

Blackout windows were configured on each appliance to ensure that replication tasks would only start on our request. On completion of both full backups, the blackout window was removed allowing both remote appliances to replicate simultaneously to the head office appliance.

On completion, a data rate change of 10% in 10% of files was carried out on each remote office appliance and then local incremental backups taken. The graphs below show the times taken for replication runs after both full and incremental backups and compared with the theoretical times these jobs would have taken without LBR.

Many-to-One Replication Test

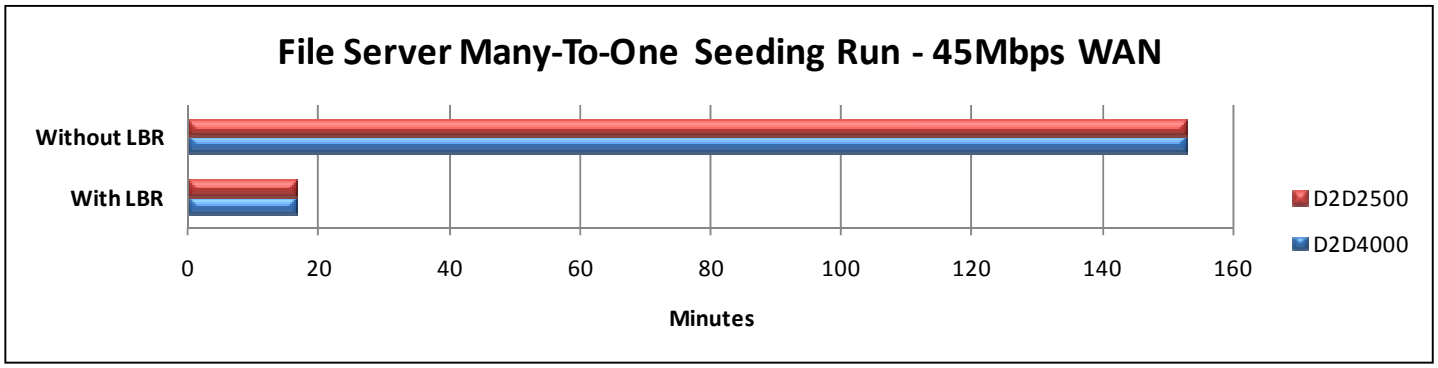

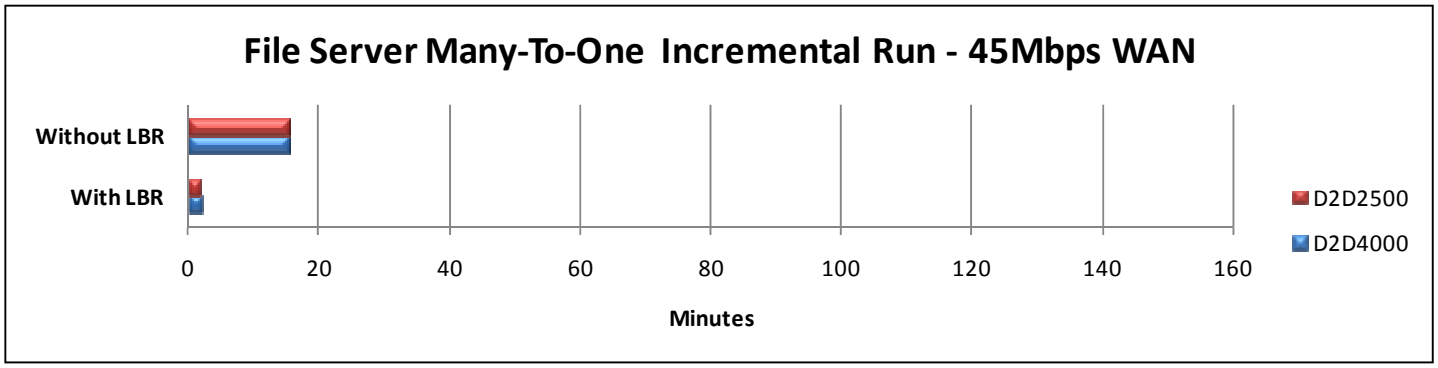

The performance results from our Many-To-One tests show that LBR offers considerable performance benefits when simultaneously replicating multiple source appliances to a single target. For the initial seeding runs after the first full backups had been taken we saw both tasks with LBR taking just over 16 minutes each to complete. This shows clearly the benefits of HP's data reduction and compression technologies as without LBR both tasks would have each taken over 2.5 hours.

Good time savings were also made when replicating the incremental runs as with LBR both replication tasks took less than 2 minutes. Without LBR both replications would each have taken over 15 minutes.

HP's Many-To-One configuration of the D2D appliances allows remote offices to backup their data locally to take advantage of the dynamic deduplication and to have automated replication of the backups over a WAN link to a central IT facility.

This provides essential off-site storage facilities and using the Tape Offload feature at the central appliance allows the remote site data held on it to be further archived to tape for long-term storage. Should a disaster occur at a remote site, its data is now stored not only at the head office but also in secure storage ensuring it can be easily reinstated onto new servers as part of a disaster recovery plan.

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.

The report concludes that HP's LBR technology delivers extensive performance improvements across the board for all data replication operations. The most significant findings are the huge reductions in the time taken to complete daily replication tasks. This has far reaching consequences as LBR now makes it possible for many businesses to run these operations overnight whereas with standard replication solutions this would not have been feasible.

The biggest time savings were demonstrated over a 2Mbps WAN link. With the 50GB data sample, the full one month of SQL Server replication tests with a 5% daily data change rate were all completed in 27 hours. If the test data had been replicated to the remote site with no data reduction applied it would have taken 395 hours representing a time saving for one month of 368 hours – or 15.4 days.

When run over a 2Mbps WAN link, the tests for Exchange, File Server and the SQL Server with a 2% daily data rate change all delivered time savings for the one month test period of between 11.6 and 13.2 days. For the tests over a 10Mbps WAN link we saw replication time reductions over one month of between 55-73 hours and for the 45Mbps tests these returned time savings of 11-14 hours.

The report shows that the greatest time savings are made with the slowest WAN links. The same fixed 50GB data sample was used in all the tests and therefore takes longer to replicate over lower bandwidth links so the same percentage changes applied to the test data returns a larger time saving.

The results also show that a wide range of data types can take advantage of HP's deduplication and LBR technologies. However, users will see variations in time savings according to the content being backed up and replicated. Files such as audio, video and graphics that have already had compression applied result in the least gains whereas common files such as documents and spreadsheets will see the most.

A key finding is that the performance improvements demonstrated by HP's LBR will allow businesses to run daily replication tasks over lower bandwidth WAN links than was previously thought possible. With high speed WAN links representing a significant outlay the cost savings will be considerable.

When calculating the time savings the initial seeding runs were not included as the majority of businesses will conduct these locally over a fast LAN connection or use the tape offload functions. However, our initial runs show it is feasible to use a low cost 2Mbps WAN link to seed a remote appliance with 50GB of backup data inside a weekend.

Compliancy with data protection regulations calls for secure, off-site storage of critical data for disaster recovery but traditional solutions have made this a costly and time consuming exercise for SMEs. HP's D2D4000 appliances now offer these very businesses a cost-effective solution for data replication and combining its Dynamic Deduplication and Low Bandwidth Replication into a single platform delivers significant benefits.

Using a single management GUI for all operations reduces the time required for deployment and daily maintenance and support for a wide range of replication scenarios allows business to easily implement automated backup from multiple remote sites to the head office.

The performance improvements demonstrated by LBR introduce a greater flexibility with replication operations and the decreased run times ensure that overnight windows can be met without having a negative impact on daily operations. Add in the fact that LBR enables replication to be run over lower cost, low bandwidth WAN links and the D2D products looks an ideal solution for cost-sensitive SMEs.

HP Publication number 4AA2-9115ENW

Full Test Results 2Mbps WAN

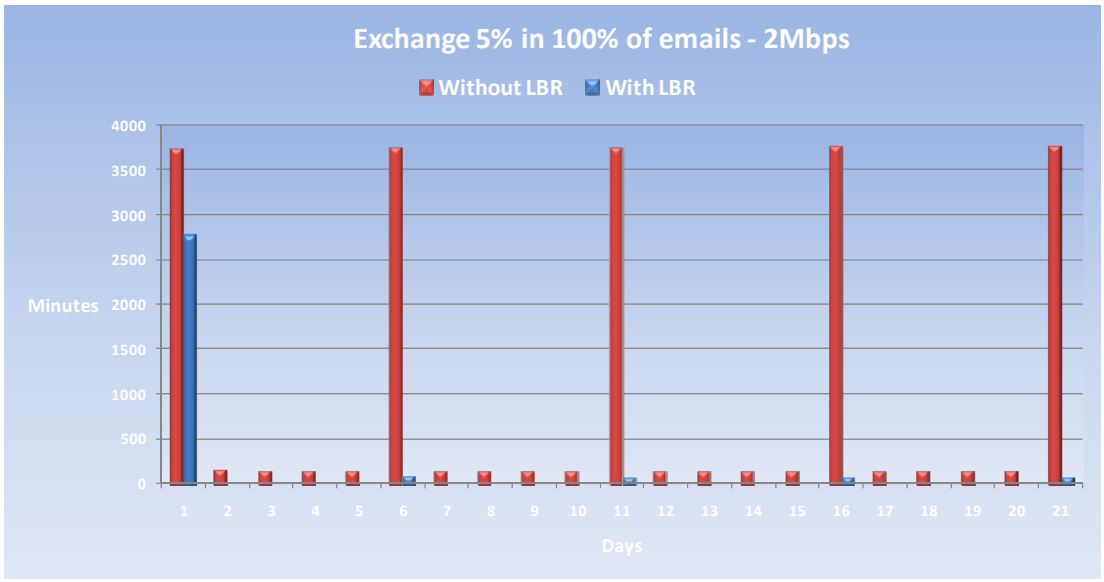

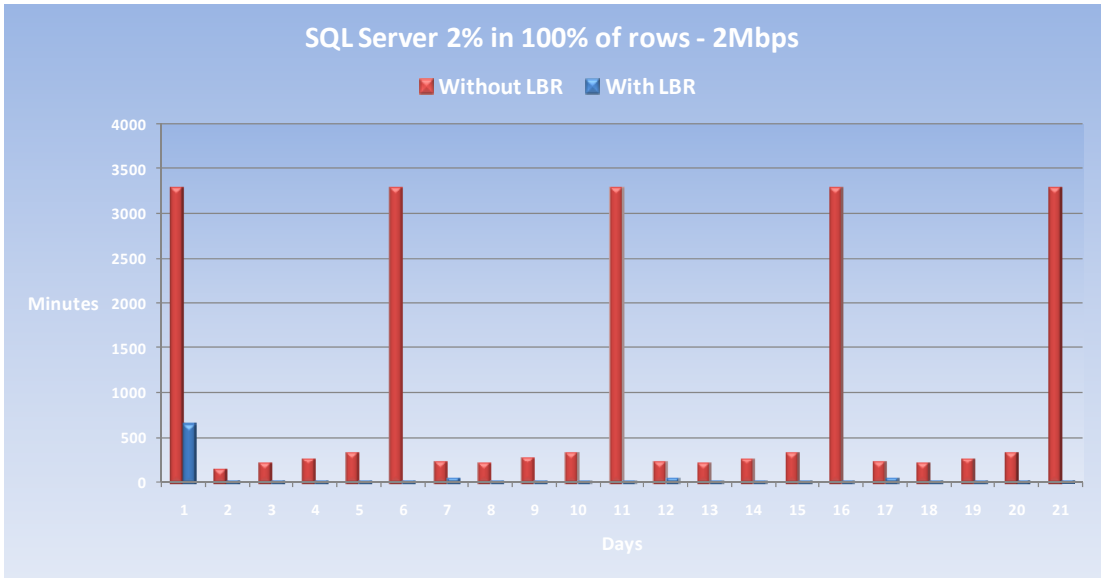

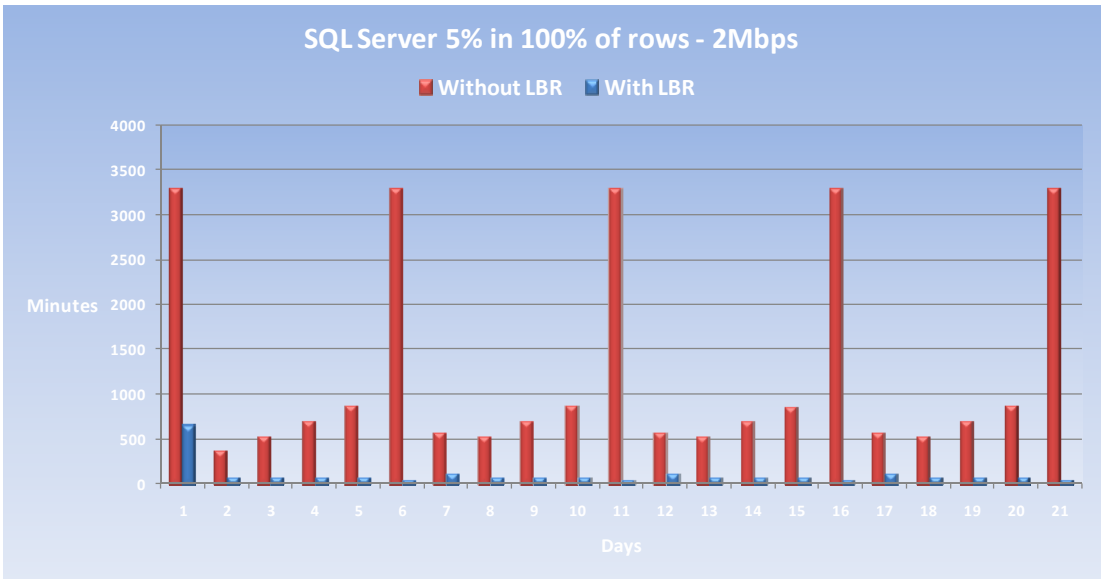

Full Test Results 2Mbps WAN

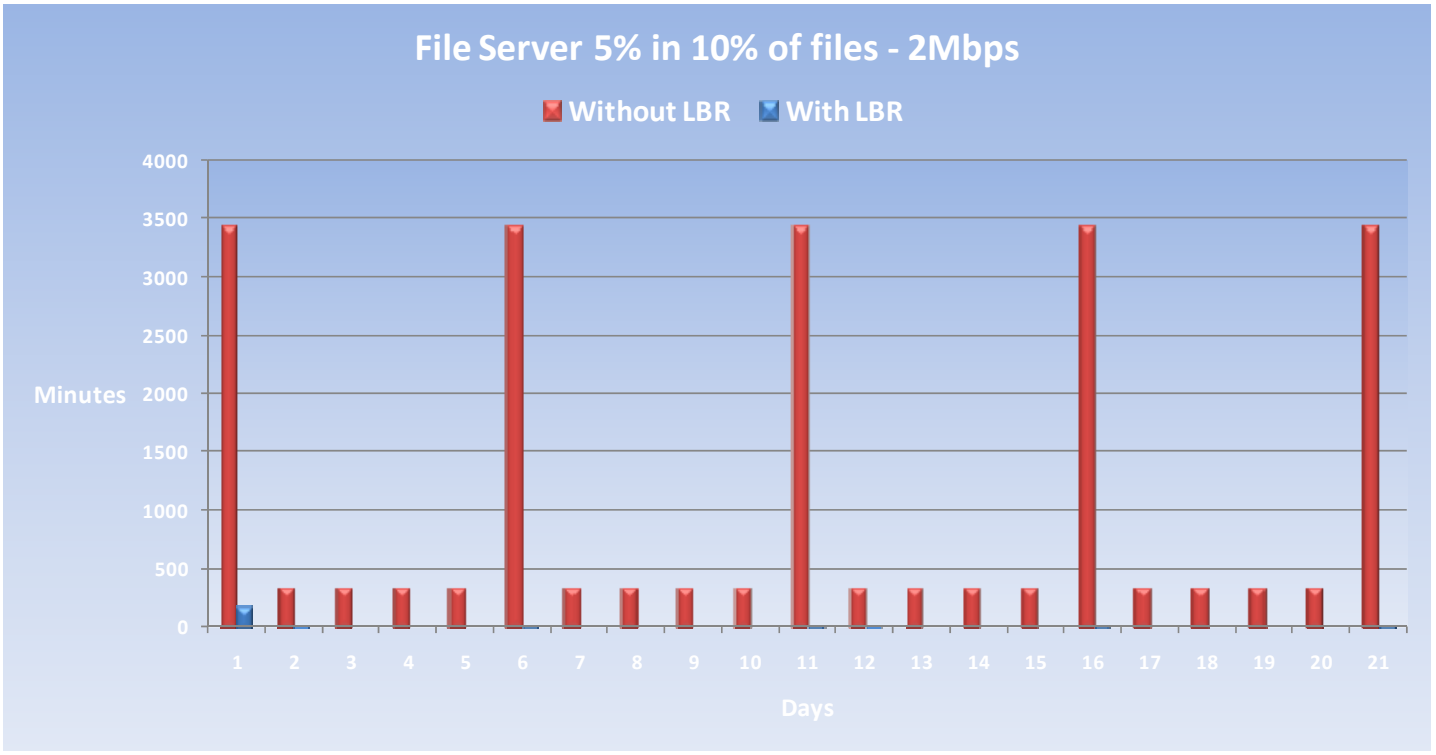

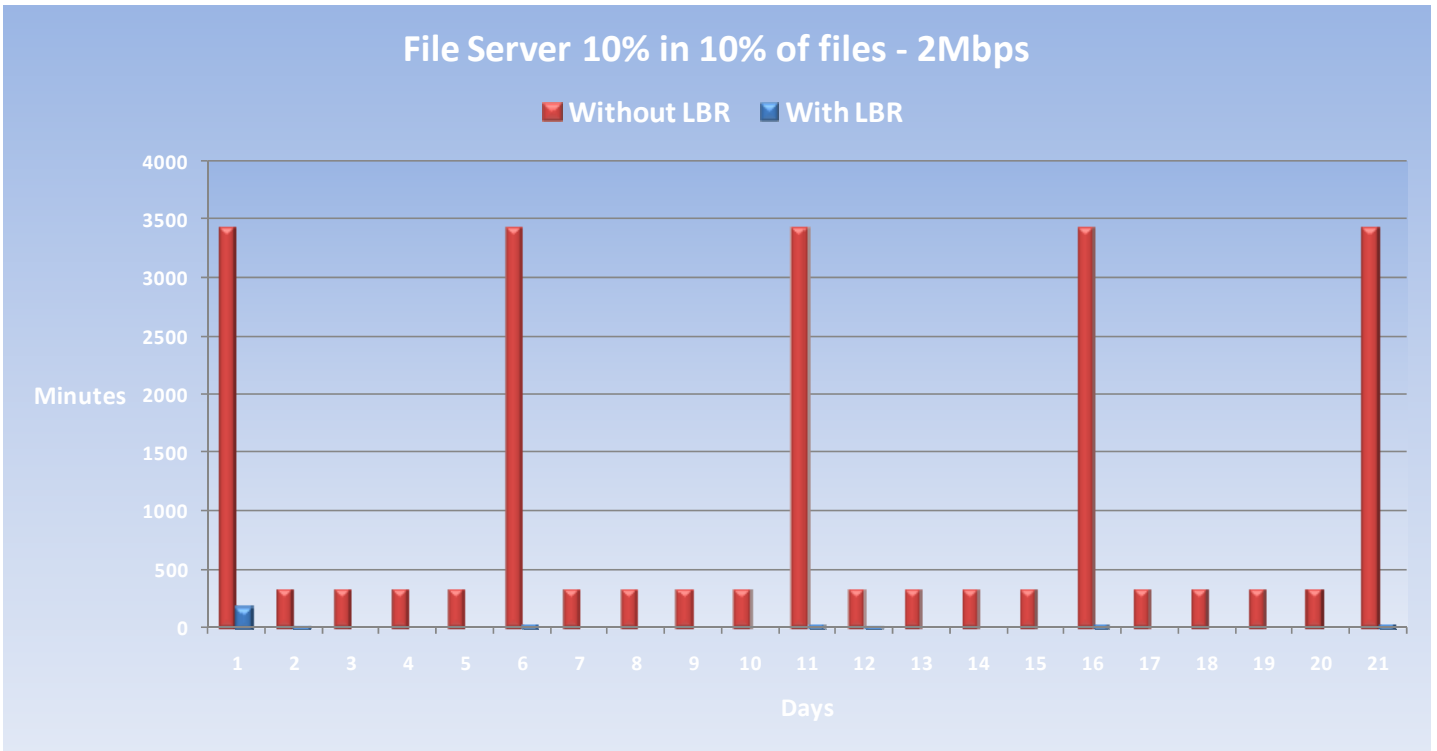

Full Test Results 10Mbps WAN

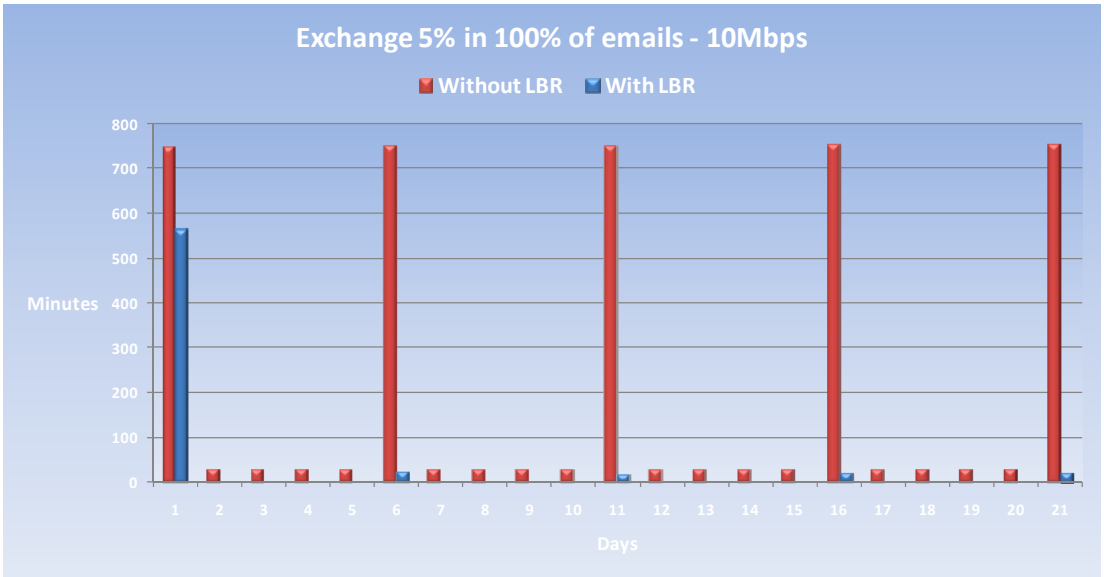

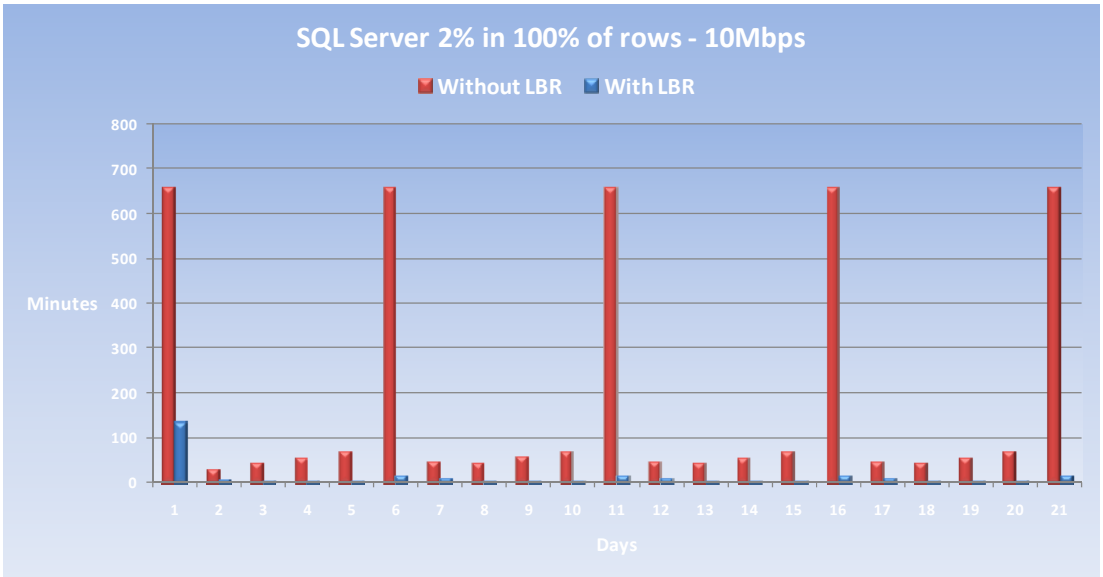

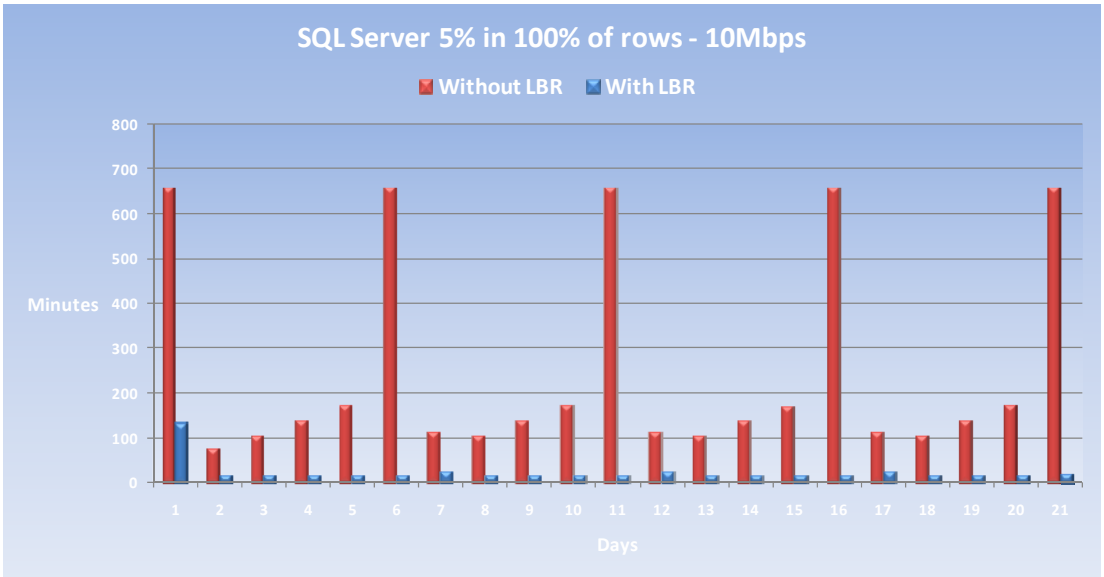

Full Test Results 10Mbps WAN

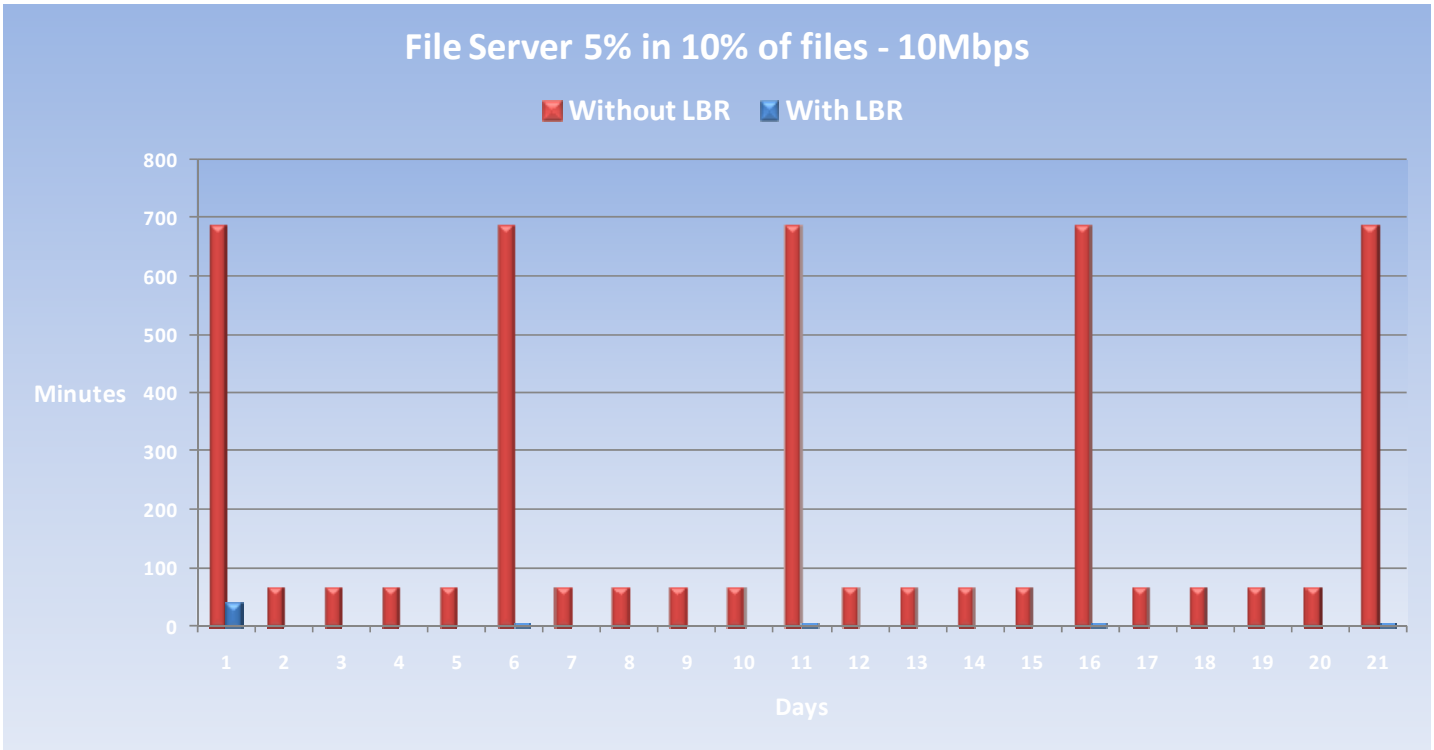

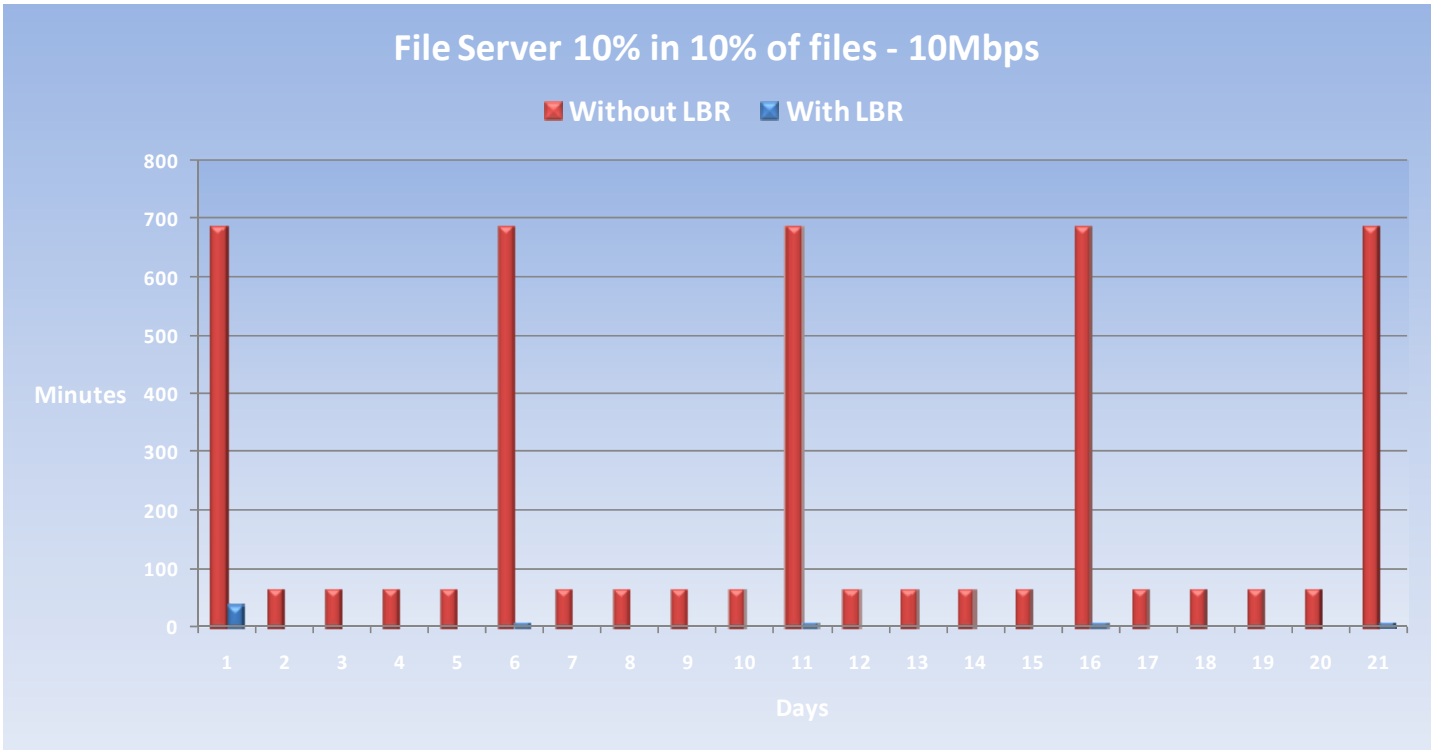

Full Test Results 45Mbps WAN

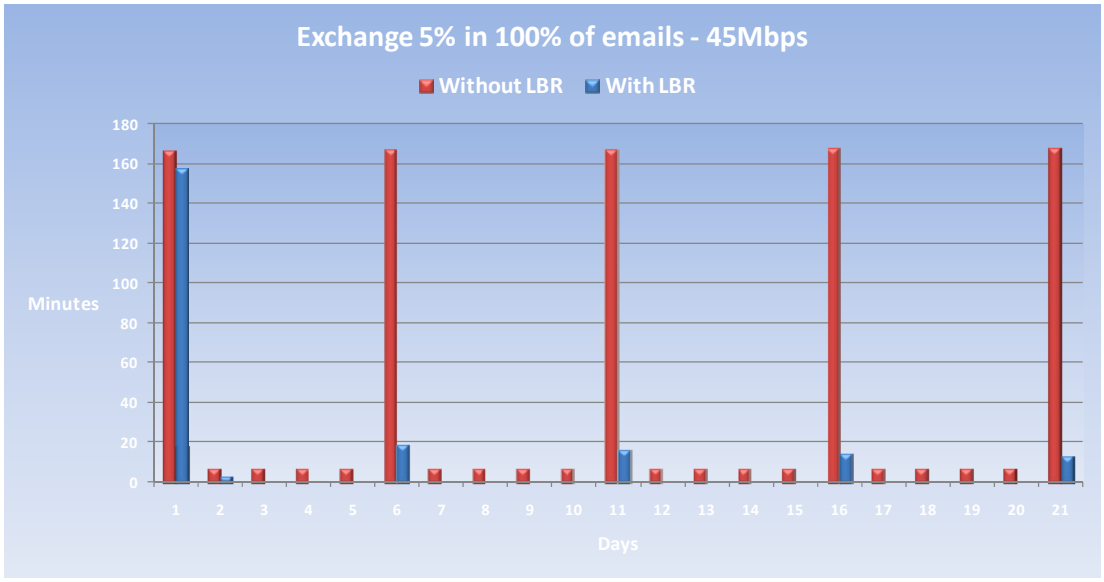

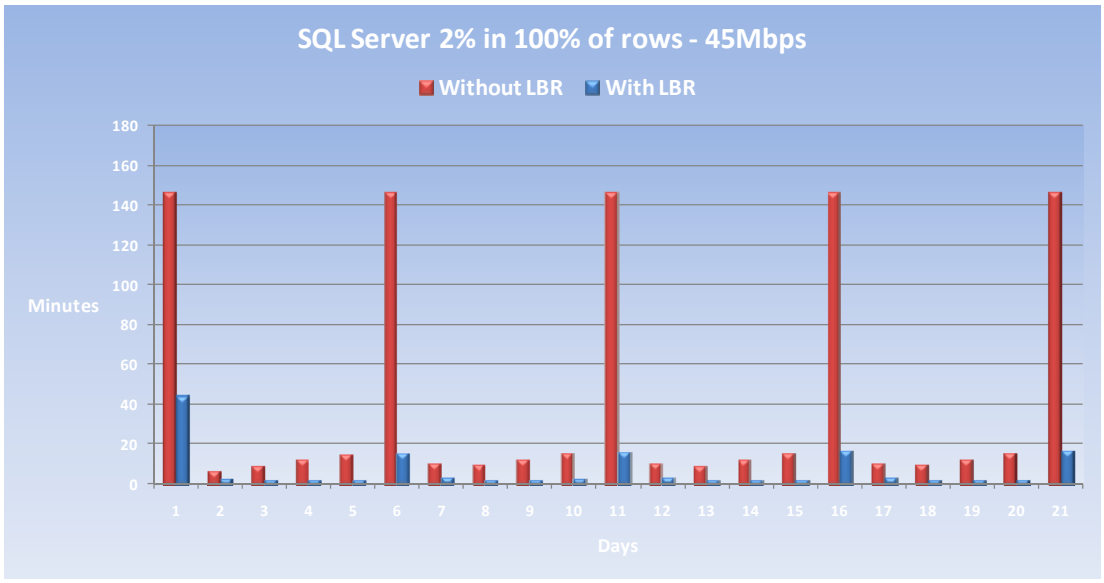

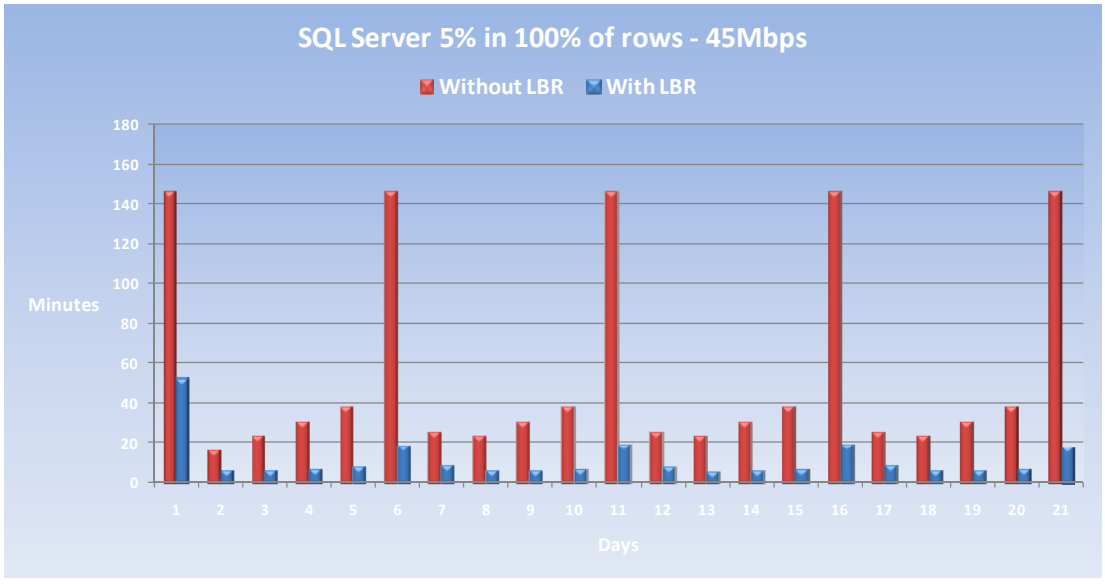

Test Results 45Mbps WAN