

DEFEND

Data Protection for businesses with remote offices across multiple locations

Business white paper

In today's information age, protecting critical data of an organization's branch offices, across multiple locations, should be standard practice. However, delivering a robust data protection strategy in environments with limited resources and often untrained personnel, can prove challenging. IDC quantified the size of this challenge: one-fifth of large companies have over 50% of their data in remote offices¹, and another one-third of large companies have 20-50% of data in their remote offices. Very often, business-critical data at remote or branch office (ROBO) locations is inadequately protected, exposing the business to greater risk of lost data and lost productivity.

¹ IDC Special Study – Hyper consolidation in the Remote Branch Driving Spend - October 2009





Table of contents

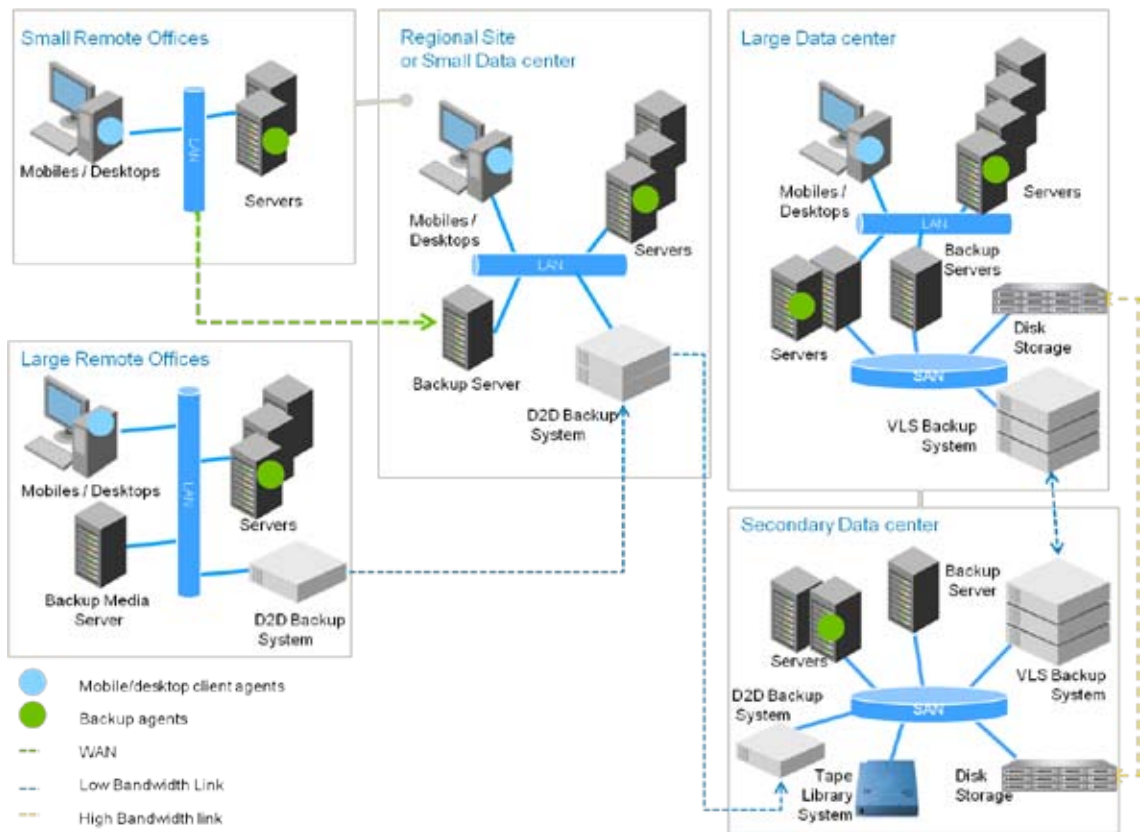
Introduction	3
Day-to-day realities for ROBO data protection	4
What is your vision of an ideal data protection solution?	4
Solution scenario summary	5
Solution scenarios	6
Scenario one: Cost-effective remote office backup with low WAN traffic	6
Benefits.....	6
Scenario two: Backup consolidation and cost-effective disaster recovery for larger ROBOs.....	7
Benefits.....	8
Scenario three: Scalability, consolidation and data center resilience	9
Benefits.....	10
HP Data Protection portfolio	11
Appendix A – HP StoreOnce deduplication is the key enabler for cost-effective data replication	12
Appendix B – HP Data protector synthetic full backup	14
Other HP Data protector technologies explained:	15
Glossary of terms	16
For more information	17
Call to action	17

Introduction

A study conducted by industry analyst Enterprise Strategy Group found that the top three IT priorities for remote office / branch office (ROBO) locations were driven by business priorities: to improve information security, ensure regulatory compliance, and enhance disaster recovery². However, the many choices in data protection technologies and approaches, coupled with a wide range of vendors to choose from, can make data protection planning a daunting endeavor.

This white paper drills deeper into these challenges and the considerations for a better way to approach data protection. It also explains how the HP data protection portfolio can help businesses overcome their data protection challenges, and drive better business outcomes. The white paper includes a glossary, defining terms used to expand data protection approaches, and where customers can find further information on HP data protection offerings.

Figure 1. Enterprise deployment with small and large ROBOs



² Enterprise Strategy Group Research Report: Branch Office Optimization, 2007

Day-to-day realities for ROBO data protection

Remote office/branch office data backup and restore has always been a challenge for IT managers. Some of the challenges include:

- Facing the choice of using untrained personnel to manage the increasingly complex physical tape infrastructure, or hiring resources to manage data backup tasks at remote sites
- Paying service providers to collect, store, and retrieve tapes between remote and central offices
- Meeting a backup window without extending into normal business hours, and affecting critical business applications
- Backing up data over the company WAN to a central site being restricted because of the time and cost involved in sending relatively high volumes of data to be backed up over a relatively low-speed link
- Unable to prove to auditors that a comprehensive site disaster recovery plan exists

What is your vision of an ideal data protection solution?

A business with multiple locations should not struggle with inefficient or ineffective data protection processes. Hardware and software solutions for data backup and restore should be easy to integrate into the existing environment, provide easy-to-manage data protection regardless of the IT expertise, and clearly demonstrate cost efficiencies. This vision must also realize that the ROBO and the associated data centers are very different in terms of costs requirements and performance expectations, and require different technologies. Ideally, a data protection infrastructure should result in:

- Spending less time on physical tape management, and a reduction in the risk of human media-related errors
- Automating processes with a single management console across the entire backup environment
- Effective storage capacity utilization, retaining backup data close to hand for as long as possible for rapid single file restore
- Using backup and restore technologies for fast backup of file servers and a reduced backup window
- Easy maintenance and management of backup copies in multiple locations
- Making the best use of existing WAN links from ROBO sites to regional data centers
- The capability of multiple layers of data protection to meet with ranging service level agreements, and levels of risk.

With HP ROBO solutions, data is accessible wherever it is required; in some cases, this may be in three different locations and on physical tape for greater flexibility in restore or disaster recovery (DR) options. With ability to manage data backup or data restore through a single backup application, IT personnel with relevant expertise can recover the data from the ROBO appliance at multiple ROBO locations, and manage the data recovery process from the main/regional data center.

HP StorageWorks D2D (Disk-to-Disk) Backup Systems, HP StoreOnce deduplication technology and HP Data Protector software can provide a complete solution for distributed office environments. Additionally, HP StorageWorks VLS (Virtual Library Systems) provide higher capacity, higher performance deduplication solutions for enterprise data centers.



Solution scenario summary

For an overview of typical profiles, data protection issues, and HP data protection capabilities, this table provides a quick comparison across three typical scenarios. Detailed scenarios are visually represented in the next section.

Table 1. Solution scenario summary

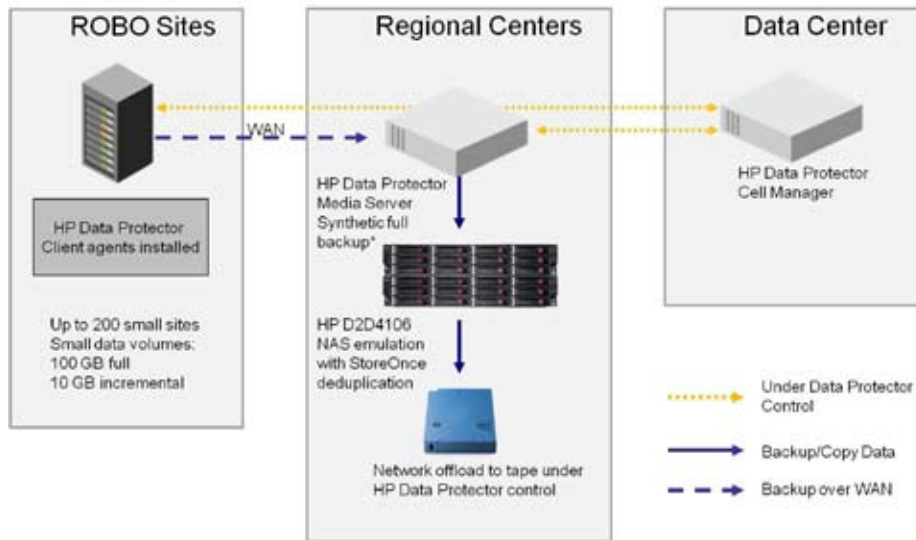
	Scenario one	Scenario two	Scenario three
Typical sizing	200+ ROBO's with full backup size of 100 GB and incremental backups of 10%	100+ ROBO's with 200-500 GB full backups	100+ ROBO's with 200-1 TB full backups PLUS Up to 4 regional data centers with 7.5TB to 300TB backups requiring DR solution
Problem	Large number of ROBO sites with small backup volumes	Significant backup volume at remote sites and a requirement for local recovery at the ROBO site	All above data requiring to be consolidated at main data centre
Solution	HP Data Protector Synthetic full backups over WAN to HP D2D with NAS Emulation	<ul style="list-style-type: none"> Up to 50:1 fan-in from small HP D2D units in the ROBO to larger consolidated single HP D2D4312 in the data centre HP Replication Manager software to manage multi-site replications. Offload to physical tape if required 	<ul style="list-style-type: none"> HP D2D for ROBO's, replicating to a larger D2D at the Regional Data Center HP D2D or VLS copies of local data at the regional data centre using HP data protector copy functionality Additional replication to D2D or VLS in the main data center for disaster recovery
Why HP	<p>Synthetic full functionality allows object consolidation and ensures only small volumes of data transmitted over WAN</p> <p>Scalable up to 100s of ROBO sites</p> <p>Recovery to ROBO from data centre under data centre control</p>	<p>Only HP offers a low cost entry point to deduplication replication for ROBO's (sub \$5000)</p> <p>Unique HP StoreOnce deduplication technology</p> <p>Target D2D can be periodically tape to tape copied to physical tape using special functions in HP Data Protector</p> <p>End to End control from HP data protector</p>	<p>Only HP offers both a low cost entry point to deduplication and replication (using D2D) AND a choice of highly scalable, appliances (D2D or VLS) for data center environments</p> <p>Target VLS can be periodically tape to tape copied to physical tape using special functions in HP data protector</p> <p>End to End control from HP data protector</p>

Solution scenarios

The scenarios discussed below cover a wide range of implementations that are possible using HP D2D Backup Systems and HP Data Protector Software, from a small business that has several outlets to a major corporation that has hundreds of remote/branch offices, several regional data centers and multiple main data centers.

For information on the enabling technologies for these solutions, please refer to the appendices.

Figure 2. Client based backup over the WAN using HP Data Protector Synthetic Full functionality



Scenario one:

Cost-effective remote office backup with low WAN traffic

Solution:

Client-based backup over the WAN using HP Data Protector synthetic full functionality

In this first scenario, we follow a basic client agent backup process, suitable for a large number of ROBO sites where relatively small amounts of data need to be protected. In order to reduce the amount of data that needs to be transmitted over the WAN, customers can use the “synthetic full” feature of HP Data Protector. After the initial full backup only incremental backups need to be performed. The backup data is sent to a “File Library” device (in this case a NAS share with deduplication on the D2D4106 unit), where a “synthetic full” can be created at any time for a restore. Alternatively, there is another feature in HP Data Protector called “virtual full” that is similar to synthetic full, but uses less storage space. However, it can only be implemented on a file library device, such as an HP P2000 disk array. Finally, under HP Data Protector control, the synthetic fulls can be copied to physical tape for long-term archiving.

Note: the HP StorageWorks D2D4106 Backup System offers up to 9TB of usable capacity. This can be extended to up to 180TB over time with HP StoreOnce deduplication at a factor of 20x.

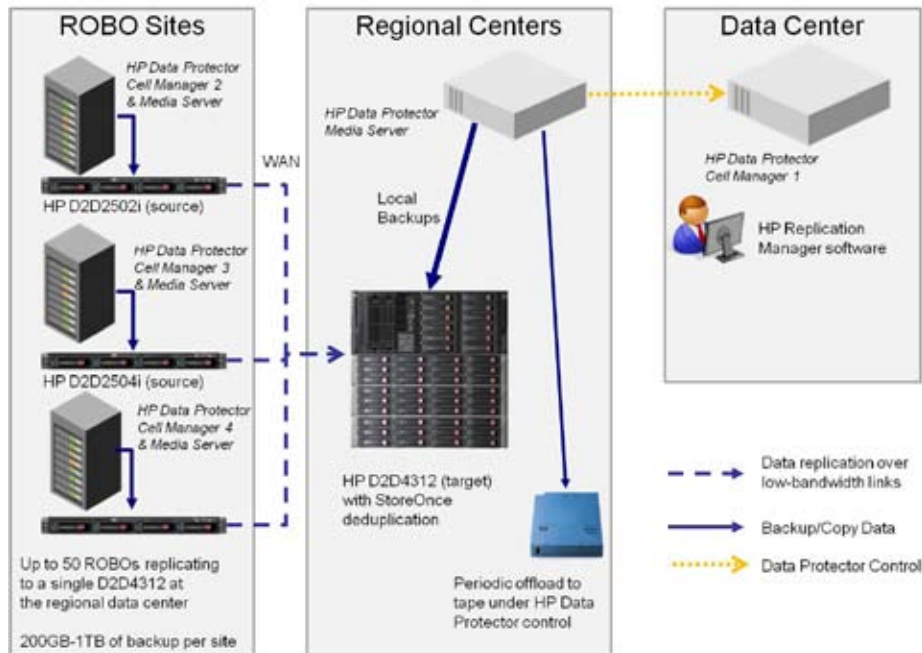
Benefits

- Object consolidation using synthetic full backups to reduce WAN traffic
- Variable Retention policies. Short term on D2D, long term on tape for archive
- Fast recovery from D2D4106 with NAS emulation
- HP Data Protector manages the entire process from the main data centre
- Reduced WAN utilization. Not in this case, it is synthetic full that reduces traffic not D2D

However, there are some constraints associated with this solution, including:

- It does not provide a local source of the data for recovery. A synthetic full created at the regional centre is required and the necessary files to be restored need to be transferred over the WAN.
- The time for this recovery process to complete depends on the size of the WAN connection and the amount of data being transmitted.

Figure 3. Large scale ROBO deployment - end-to-end replication with regional data center consolidation



Scenario two:
Backup consolidation and cost-effective disaster recovery for larger ROBOs

Solution:
HP StoreOnce enabled replication using HP StorageWorks D2D2502/4i in ROBO and HP StorageWorks D2D4312 in Regional Center all managed by HP Data Protector - for up to 50 remote sites.

In this scenario, the volumes of data on the remote sites are much larger, around 500 GB, and customers really need to have local restore capability. Each remote site has a local D2D2500 series appliance (D2D2502 or D2D2504i) with HP StoreOnce deduplication and the D2D uses the virtual tape library emulation mode. A user capacity of 1.5 and 3TB can yield between 30 and 60TB of backup storage with a 20:1 deduplication ratio after several backups have been retained over two – three months. If long retention is not required then the D2D2500 can support even larger backups retained for shorter periods.

Direct restore at the ROBO from the D2D unit is possible since each ROBO site contains a complete Data Protector environment of cell manager and media server and is “self-sufficient” from a restore perspective. The local D2D2500s (sources) can also be configured to replicate only the changed data across a low-bandwidth link to the regional center where a larger D2D4312 (target) device resides. A D2D4312 can support a fan-in of up to 50 D2D sources and only requires a single HP Replication Manager software license for target devices to enable replication. Source devices use Replication Manager at no charge. This allows major storage consolidation benefits at the regional data center. The replication process replicates the changes in virtual cartridges from the ROBO to the regional center.

The HP D2D replication license also entitles the user to HP Replication Manager software to provide central management capabilities for multi-site replication.

The HP Data Protector Cell Manager can be scripted to poll the D2D4312 on a regular basis to detect replicated cartridges and update its internal database (IDB) with their contents. As a result, in the event of a disaster at the ROBO, the regional center is immediately aware of which virtual tapes are required for recovery. Finally, data from the D2D4312 at the regional center can be easily copied to physical tape at regular intervals using the “object copy” functionality within HP Data Protector. In addition to acting as a replication target for ROBO sites, the D2D4312 can also act as a source for backups from the regional center servers as seen in Figure 3. The D2D4312 offers both dual 8Gb FC and dual 10 GbE connections for the local backups in the Regional center.

This solution provides robust disaster recovery. If a disaster occurs at the source (ROBO), the data remains safe at the target (regional data center). The ROBO site can then be re-constructed in any of the following ways:

- a. Re-building a new ROBO server and restoring data at the regional center using the replicated data
- b. Re-building the server at the ROBO site and reverse replicating the critical data from the D2D4312 in the regional center
- c. Offloading critical data to tape at the regional center and transporting it back to the ROBO with the necessary tape hardware for recovery

The new area of interest here is the low-bandwidth link. For the majority of ROBO sites link speeds as low as 2 MB/sec can be used (dependent on volumes of data to be replicated). Different usage models may be used to enhance the business benefits of network efficient replication, including:

- Use 100% of the link bandwidth to complete replication in a fixed window (12 hours). This is for businesses that are not necessarily 24x7 and can afford a window dedicated to replication.
- Within the D2D unit it is possible to configure “throttling” of the link and allow the replication to take place over a 24 hour period. This is for businesses that also use the WAN link for applications, and where giving the whole link over to replication would affect other application performance.

To calculate the size of the regional center D2D device and the size of the link required (either with fixed window or with throttling) HP provides a comprehensive sizing tool available at:

<http://www.hp.com/go/storageworks/sizer>

Benefits

- Locally performed and controlled backups, reliability of backups not dictated by WAN availability, and remote media server as in scenario one
- Fast local recovery when needed from two backup copies (one local and one remote)
- Ability to store several months of backups locally on the D2D at the ROBO using HP StoreOnce deduplication
- Significant bandwidth reduction—up to 95% savings in bandwidth compared to all contents being replicated—rather than just change data detected by deduplication
- Cost-effective consolidation in the regional center up to 50 ROBO D2Ds (fan-in). In addition, the D2D4312 can also be used to consolidate local regional server backup
- Highly flexible disaster recovery options
- Easy offload to tape with HP Data Protector object copy
- Management of the entire process from the main data center by HP Data Protector
- Option to manage single ROBOs from the regional/main data center

As for scalability – the example above shows a 50:1 fan-in, if a customer environment has more than 50 ROBO sites, the customer would need multiple D2D4312 devices at the regional data centre.



Scenario three:
Scalability, consolidation and data center resilience

Solution:
Network-efficient replication over low-bandwidth links using HP StorageWorks D2D for replication between ROBO sites and regional centers, and HP StorageWorks VLS Backup Systems with HP Data Protector software for replication between regional centers and the main data center. HP Replication Manager software facilitates multi-site replication through central management capabilities.

In this scenario, the volumes of data involved at the regional data centers are significantly high and the need for scalability paramount. HP enterprise-level Virtual Tape Libraries (VLS) are used because they can scale higher in terms of capacity and performance than the HP D2D Backup Systems. An alternative configuration could be multiple HP D2D4312 Backup Systems.

The ROBO to regional center solution is the same as in scenario two. However, instead of local regional data center backups being done to the HP D2D, the regional data center volumes (greater than 30 TB) are now backed up to a local VLS9000 device. In this case the VLS9000 can scale to capacities between 10TB and 1280TB of data.

VLS systems also support network efficient replication of data over lower bandwidth links, and currently support up to a 4-to-1 fan-in that enables consolidation of replications from several regional data centers in a single VLS9000 target device at the main data center. The VLS9000 uses a multi-node architecture to provide scalability in terms of backup performance, as well as deduplication and replication performance. Given the higher volumes of replicated data involved, the lower bandwidth links for VLS are typically 100 Mbits/sec upwards.

This solution with VLS allows consolidation at the main data center. The VLS9000 there is used both as a replication target for the four regional data centers and for main data center backups as well, providing an excellent solution for consolidation.

Replicated cartridges from the regional center arriving at the VLS target in the main data center send an e-mail on an hourly basis to the Data Protector Cell Manager at the main data center. The Data Protector scripts then import these newly replicated tapes into the backup catalog at the main data center to allow easy access to replicated data when disaster recovery becomes necessary for the regional data center.

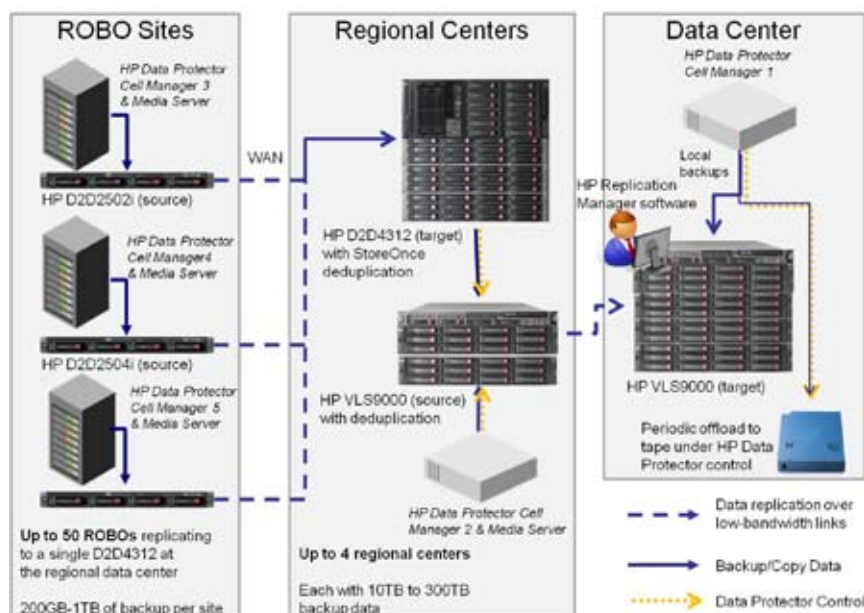
A further enhancement to this scenario would be to have active-active replication, where the main data center backups to the VLS are also replicated back to the regional data center, giving the main data center a very cost effective disaster recovery process using the existing infrastructure. Due to the different nature of D2D and VLS deduplication technologies, it is not possible for D2D units to replicate directly into a VLS unit.

The disaster recovery options for the regional data center using VLS are identical to those used for D2D. Again, with the higher volumes of data involved between regional data centers and main data centers, HP expects the usage model for VLS to VLS replication to be predominantly:

- A dedicated replication link that is not shared with application traffic
- 24-hour replication window available

VLS replication sizing is also fully supported in the HP StorageWorks Sizing Tool at: <http://www.hp.com/go/storageworks/sizer>

Figure 4. ROBO and regional data to primary data centers – flexible, scalable, and resilient





Benefits

- Fast local recovery when needed at ROBOs, regional data centers, and main data centers
- Option for ROBO data to be copied to the VLS in the regional data center for onward replication to the main data center
- Ability to store several months of backups locally at the regional/main data center using VLS accelerated deduplication
- Cost-effective consolidation in the main data center, up to four regional VLS units (fan-in) and VLS at the main data center can be used for local and main data center backups as well
- Highly flexible disaster recovery options
- Easy offload to tape with HP Data Protector object copy
- Complete process management from the main data center with HP Data Protector
- Main data center disaster recovery capability if active-active replication is deployed in the future, using the existing infrastructure

HP Data Protection portfolio

The three scenarios described above give you an insight into the various ROBO and regional data center disaster recovery deployments that are possible using the following HP storage components:

HP D2D Backup Systems with HP StoreOnce deduplication

Provide disk-based backup, retaining up to 50x more data on disk, with low bandwidth replication for cost-effective data protection of distributed office environments with multiple servers.

HP Virtual Library Systems

Offer high-performance, scalable data protection solutions in demanding data center environments with accelerated deduplication and enabled replication.

HP Tape Systems (standalone tape drives and tape automation products)

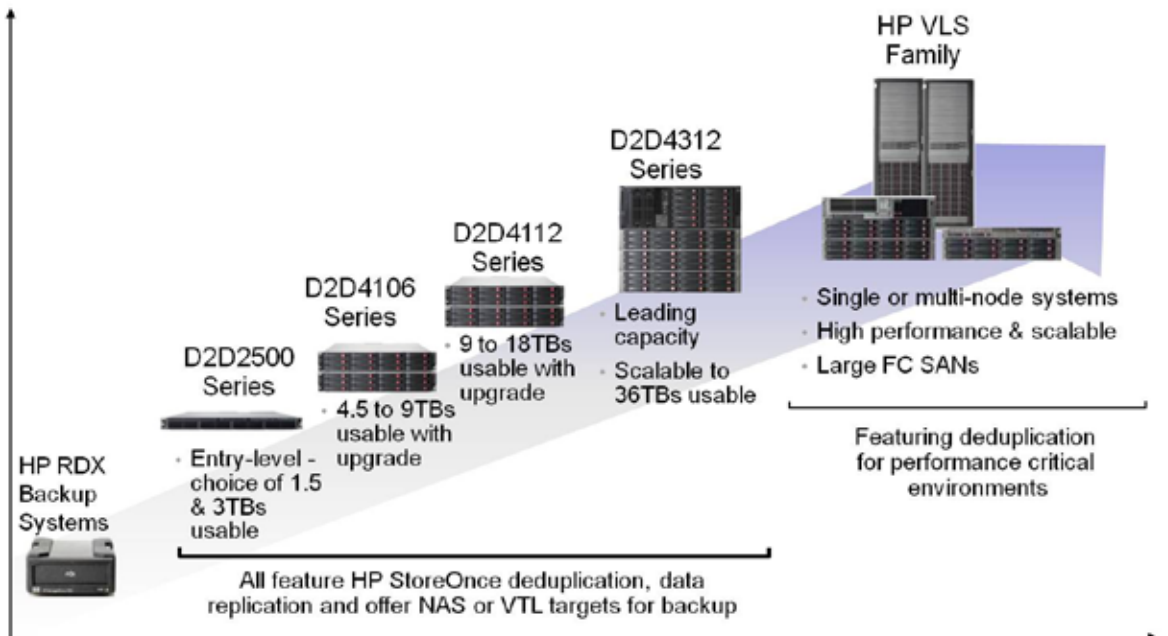
Choice of physical tape libraries designed to provide unattended backup, archive, and disaster recovery for businesses of any size.

HP Data Protector software

With limited resources and tight budgets, HP Data Protector software automates high-performance backup and recovery from disk or tape to enable 24/7 business continuity at a compelling price-to-performance ratio.

Using the wide range of components allows flexibility and scalability in the variety of ROBO and data center backup solutions that can be constructed. With innovative technologies such as the synthetic full functionality, deduplication, and low-bandwidth replication on VTL products, HP can deliver a highly effective solution that helps provide peak performance at the best possible cost for each specific location, while enabling central management and control.

Figure 5. HP disk-based data protection products



Appendix A – HP StoreOnce deduplication is the key enabler for cost-effective data replication

Data deduplication is a method of reducing storage needs by eliminating redundant data, so that over time only one unique instance of the data is actually retained on disk. Data deduplication works by examining the data stream as it arrives at the storage appliance, checking for blocks of data that are identical and eliminating redundant copies. If duplicate data is found, a pointer is established to the original set of data as opposed to actually storing the duplicate blocks—removing or “de-duplicating” the redundant blocks from the volume. However, indexing of all data is still retained so that it can be “rehydrated” should that data ever be required.

The key here is that the data deduplication is being done at the block level to remove far more redundant data than deduplication done at the file level (called single-instancing), where only duplicate files are removed.


By its very nature, the backup process tends to generate a great deal of repetitive copies of data, and hence data deduplication is especially powerful when it is applied to backup data sets. The amount of redundancy will depend on the type of data being backed up, the backup methodology and the length of time the data is retained.

Once the original file is stored the technology eliminates duplicate data down to the block or byte level on all future changes to that file: If a change is made to the original file, then data deduplication saves only the block or blocks of data actually altered, (a block is usually quite small, less than 10KB of data.) So let’s say the title of our 1 MB presentation is changed. Data deduplication would save only the new title, usually in a 4 KB data block, with pointers back to the first iteration of the file. Thus, only 4 KB of new back up data is retained.

When used in conjunction with other methods of data reduction, such as conventional data compression, data deduplication can cut data volume even further.

HP StoreOnce deduplication software delivers improved performance at half the price point of competitive solutions and enables clients to spend up to 95 percent less³ on storage capacity compared to traditional backup. HP StoreOnce software is available in all HP StorageWorks D2D backup systems, providing a single technology that can be deployed at multiple points in a converged infrastructure.

Figure 6. Illustrating the deduplication process



Appendix A - continued

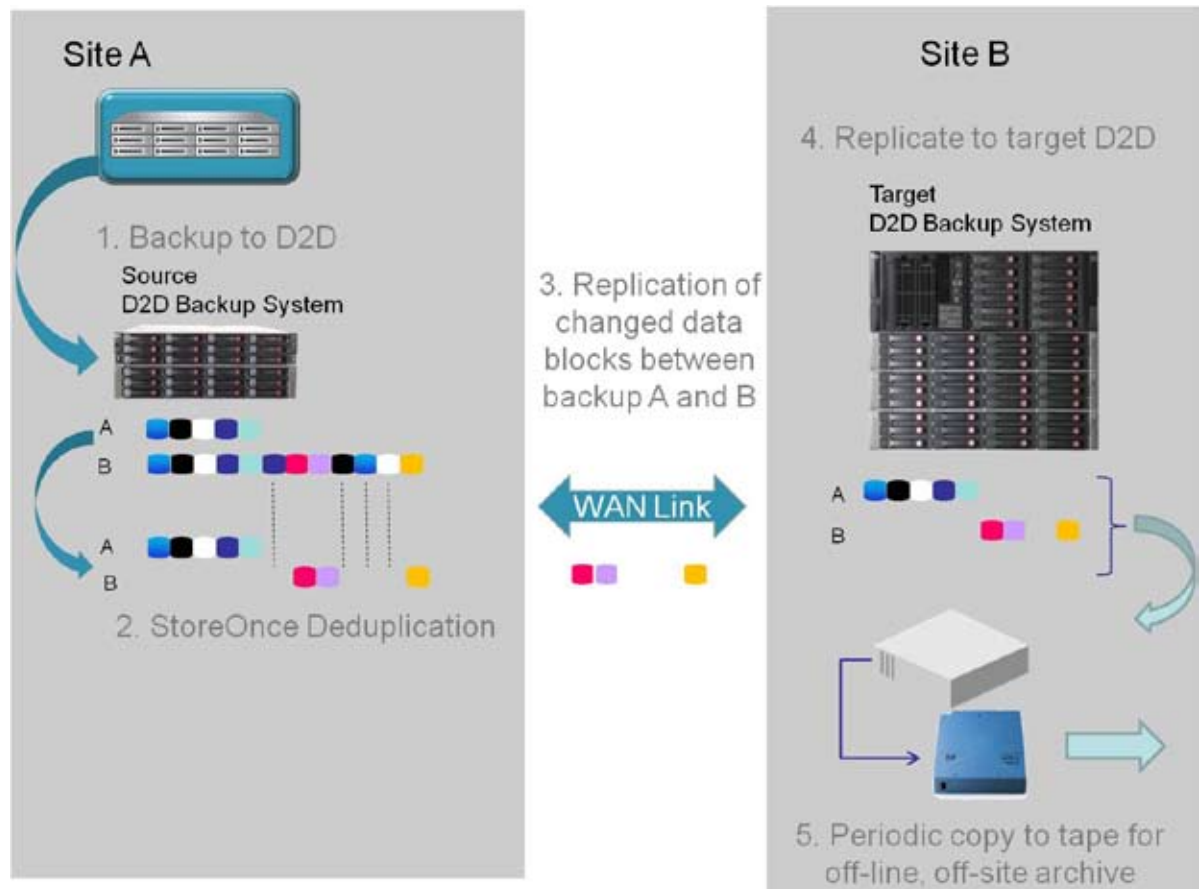
Replication of data can occur between devices within the same product family i.e. between the HP StorageWorks D2D2500, D2D4000, D2D4112 and D2D4312 Backup Systems. It is not possible to replicate data from D2D to the VLS product ranges nor to competitors' disk based backup systems with replication features. HP D2D & VLS devices use a different implementation for deduplication, each designed to be most appropriate to the target market in terms of ease of use and affordability.

Data replication is the process of making a replica copy of a data set across a network to a "target site". It is generally used to transmit backup data sets off-site to provide disaster recovery (DR) protection in the event of catastrophic data loss at the "source site". In the past, only the largest companies could afford to implement data replication as replicating large volumes of data backup over a typical WAN is expensive.

However, because StoreOnce deduplication knows what data has changed at a block or byte level, replication becomes more intelligent and transfers only the changed data as opposed to the complete data set. This saves time and replication bandwidth making it possible to replicate data over lower bandwidth links for a more cost-effective, network efficient replication solution.

HP D2D Backup Systems with StoreOnce provide an automated and practical disaster recovery solution for a wide range of data centers, in addition to an ideal solution for centralizing the backup of multiple remote offices.

Figure 6. Deduplication is key to enabling more affordable, network efficient data replication between sites. Using HP D2D meets restore performance service levels, adding periodic offload to tape adds a level of protection and meets long-term archival requirements.



³ Based on multi-site deployment of D2D4312 ROI calculations as replacement technology for pure tape backup and offsite archival (HP D2D TCO Analysis from May 2010)

⁴ Based on average deduplication ratio of 20:1. Deduplication ratio could be as high as 50:1 under optimal circumstances.

Appendix B – HP Data Protector Synthetic Full Backup

HP Data Protector software offers an advanced backup solution called synthetic full backup. This solution enables you to create synthetic full backups and virtual full backups with an operation called object consolidation allowing implementing an incremental forever backup strategy.

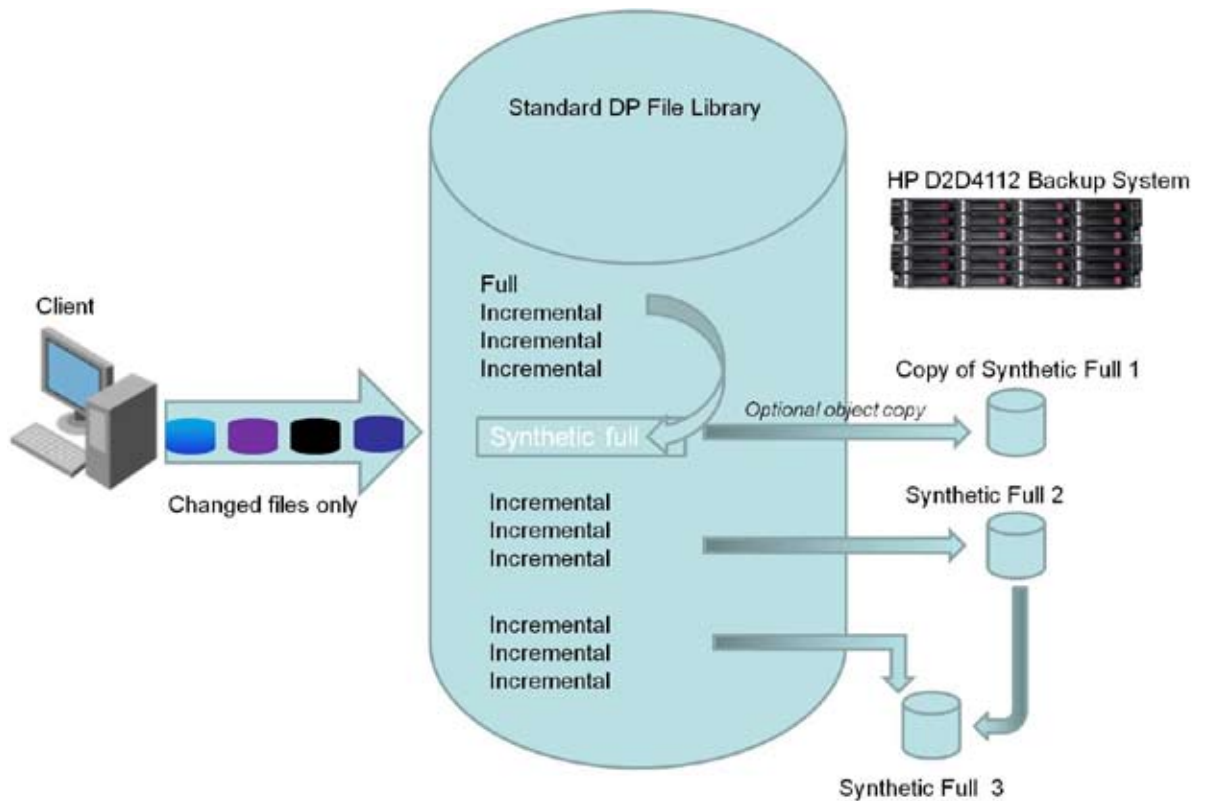


Figure 7. HP Data Protector Synthetic Full (Incremental Forever) Backup explained

Synthetic backup eliminates the need to run lengthy full backups. This technology works by merging all incremental backups into a full 'synthetic' backup—a process that can be repeated indefinitely, with no need to run a full backup again. If all the backups, full and incremental, are written to the same HP Data Protector file library which uses a distributed file media format (DFMF) e.g on an P2000 disk array, an even more efficient type of synthetic backup is possible. This is called Virtual Full Backup. The solution uses pointers to consolidate data rather than copying the data. As a result, the consolidation takes less time and avoids unnecessary duplication of data. Only the HP Data Protector Synthetic backup technique is supported on the D2DNAS emulation.

This capability also helps enable continuous data protection for Windows systems. For systems running Windows 2000 and above, HP Data Protector utilizes the built-in Microsoft Windows Change Journal to quickly generate a list of files that have been altered since the last backup was performed. This option avoids the need for timely "tree walks" to identify files that have changed, which can be especially time-consuming on systems with many small files. This allows incremental backups to run much faster and more frequently, providing near-continuous data protection at no additional cost.

Appendix B - continued

The 'incremental forever' paradigm means that except for the first backup where a full backup is performed, only incremental backups are executed. This concept presents the most efficient way of backups; only the changed data is backed up.

However, without object consolidation, the restore process would last far too long since all backup sessions would have to be restored. Due to this behavior, regular full backups are required and all backups would need to be protected permanently. Object consolidation removes this drawback.

During the object consolidation session, HP Data Protector software reads the backed up data from the source media, merges the data, and writes the consolidated version to the target media. This consolidation process taking place on the D2DNAS device does take a significant amount of time.

An object consolidation can result either in a synthetic full or a virtual full backup.

A synthetic full backup is the result of an object consolidation operation where a restore chain of backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed. HP D2D NAS emulation supports synthetic full backups.

A virtual full backup is an efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single HP Data Protector software file library using distributed file media format. This requires the File Library to be a disk array such as HP MSA2000.

Other HP Data Protector technologies explained:

DHP Data Protector Cell Manager

The Cell Manager is the main system in the cell. The Cell Manager:

- Manages the cell from a central point
- Contains the IDB (Internal Data Base) - the IDB contains information about backup details such as, backup durations, media IDs, and session IDs
- Runs core HP Data Protector software
- Runs Session Managers that start and stop backup and restore sessions and write session information to the IDB.

HP Data Protector Disk Agents

Client systems you want to back up must have the HP Data Protector Disk Agent (DA), also called Backup Agent, installed. To back up online database integrations, install the Application Agent. The Disk Agent reads or writes data from a disk on the system and sends or receives data from a Media Agent. The Disk Agent is also installed on the Cell Manager, thus allowing you to back up data on the Cell Manager, the HP Data Protector configuration, and the IDB.

HP Data Protector Media Server

Client systems with connected backup devices must have a HP Data Protector Media Agent (MA) installed. Such client systems are also called Drive Servers. A backup device can be connected to any system and not only to the Cell Manager. A Media Agent reads or writes data from or to media in the device and sends or receives data from the Disk Agent.

HP Data Protector Object Copy

The Data Protector object copy functionality enables you to copy selected object versions to a specific media set. You can select object versions from one or several backup sessions or object consolidation sessions. During the object copy session, Data Protector reads the backed up data from the source media, transfers the data, and writes it to the target media.

Glossary of Terms

HP StorageWorks Sizer Tool

This freely available tool

<http://www.hp.com/go/storageworks/sizer>

allows users to specify backup capacities, retention schemes, data change rates, replication window, link speed etc and calculates likely backup storage capacity required.

Active-Active

Replication from a Source device on Site A to a Target Device on Site B and also replication from a Sourced device at Site B to a Target Device at Site A. Other replication configurations include: Active->Passive - replication from a Source to Destination in one direction only, and Many-to-One which is replication from multiple sources to a single target device.

Deduplication ratio

The reduction in storage required for a backup (after several other backups have taken place). The ratio is highly dependent on:

- Rate of change of data (for example, 10% of the data in 10% of the files)
- Type of data—files vs database
- Retention period of backups
- Efficiency of deduplication technology implementation

Fan-In

This is the number of source appliances or libraries that can be channeled into a target appliance or library.

Link Speed

This is speed in Mbits/sec that the Telco provides point to point or Any to Any to allow replication to take place. The speeds can vary from 2 Mb/sec to 1 Gb/sec. Increasing Speed means increasing cost.

D2D Replication Licence

Only one replication licence is required per target device, this is a chargeable licence.

D2D Replication manager

This is a comprehensive utility which monitors complex D2D replication environments and is available free of charge when you redeem your D2D replication licence.

Bandwidth Utilization (Throttling)

This is the amount of available bandwidth a device can be configured to use. For example the D2D & VLS devices have network throttling so they can be configured to use for example only 20% of the available bandwidth. This prevents any single device from “hogging” the total bandwidth available, and allows replication to take place over longer periods of time without affecting application responses using the same link.

Tape Offload

The preferred method of copying data from D2D to Physical tape is to use the ISV (independent software vendor) backup software so that multiple streams can be offloaded to physical tape at once improving copy speeds. This approach also allows the backup software to track all copies of the same data through its internal database. This technique can be applied for:

- D2D VTL copies to physical tape
- D2D NAS copies to physical tape

For more information

www.hp.com/go/VLS

www.hp.com/go/D2D

www.hp.com/go/storageworks/sizer

www.hp.com/go/dataprotector

Call to action

www.hp.com/go/d2d

www.hp.com/go/robo

- Advanced Backup to Disk Integration with Virtual Tape Libraries
- Integrating HP Data Protector software with HP Data Deduplication solutions
- Disk-Assisted Backup Whitepaper
- HP Data Protector software object consolidation: "Best Practices"

Share with colleagues



Get connected

www.hp.com/go/getconnected

Get the insider view on tech trends, alerts, and HP solutions for better business outcomes

© Copyright 2009-2010 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

