



Understanding the HP Data Deduplication Strategy

Why one size doesn't fit everyone

Table of contents

Executive Summary	2
Introduction	4
A word of caution	5
Customer Benefits of Data Deduplication	7
A word of caution	9
Understanding Customer Needs for Data Deduplication	10
HP Accelerated Deduplication for the Large Enterprise Customer	11
Issues Associated with Object-Level Differencing	14
What Makes HP Accelerated Deduplication unique?	14
HP Dynamic Deduplication for Small and Medium IT Environments	16
How Dynamic Deduplication Works	17
Issues Associated with Hash-Based Chunking	19
Low-Bandwidth Replication Usage Models	21
Why HP for Deduplication?	23
Deduplication Technologies Aligned with HP Virtual Library Products	24
Summary	25
Appendix A—Glossary of Terminology	26
Appendix B—Deduplication compared to other data reduction technologies	27
For more information	28

Executive Summary

Data deduplication technology represents one of the most significant storage enhancements in recent years, promising to reshape future data protection and disaster recovery solutions. Data deduplication offers the ability to store more on a given amount of storage and replicate data using lower bandwidth links at a significantly reduced cost.

HP offers two complementary deduplication technologies that meet very different customer needs:

- Accelerated deduplication (object-level differencing) for the high-end enterprise customer who requires:
 - Fastest possible backup performance
 - Fastest restores
 - Most scalable solution possible in terms of performance and capacity
 - Multi-node low-bandwidth replication
 - High deduplication ratios
 - Wide range of replication models
- Dynamic deduplication (hash-based chunking) for the mid size enterprise and remote office customers who require:
 - Lower cost device through smaller RAM footprint and optimized disk usage
 - A fully integrated deduplication appliance with lights-out operation
 - Backup application and data type independence for maximum flexibility
 - Wide range of replication models

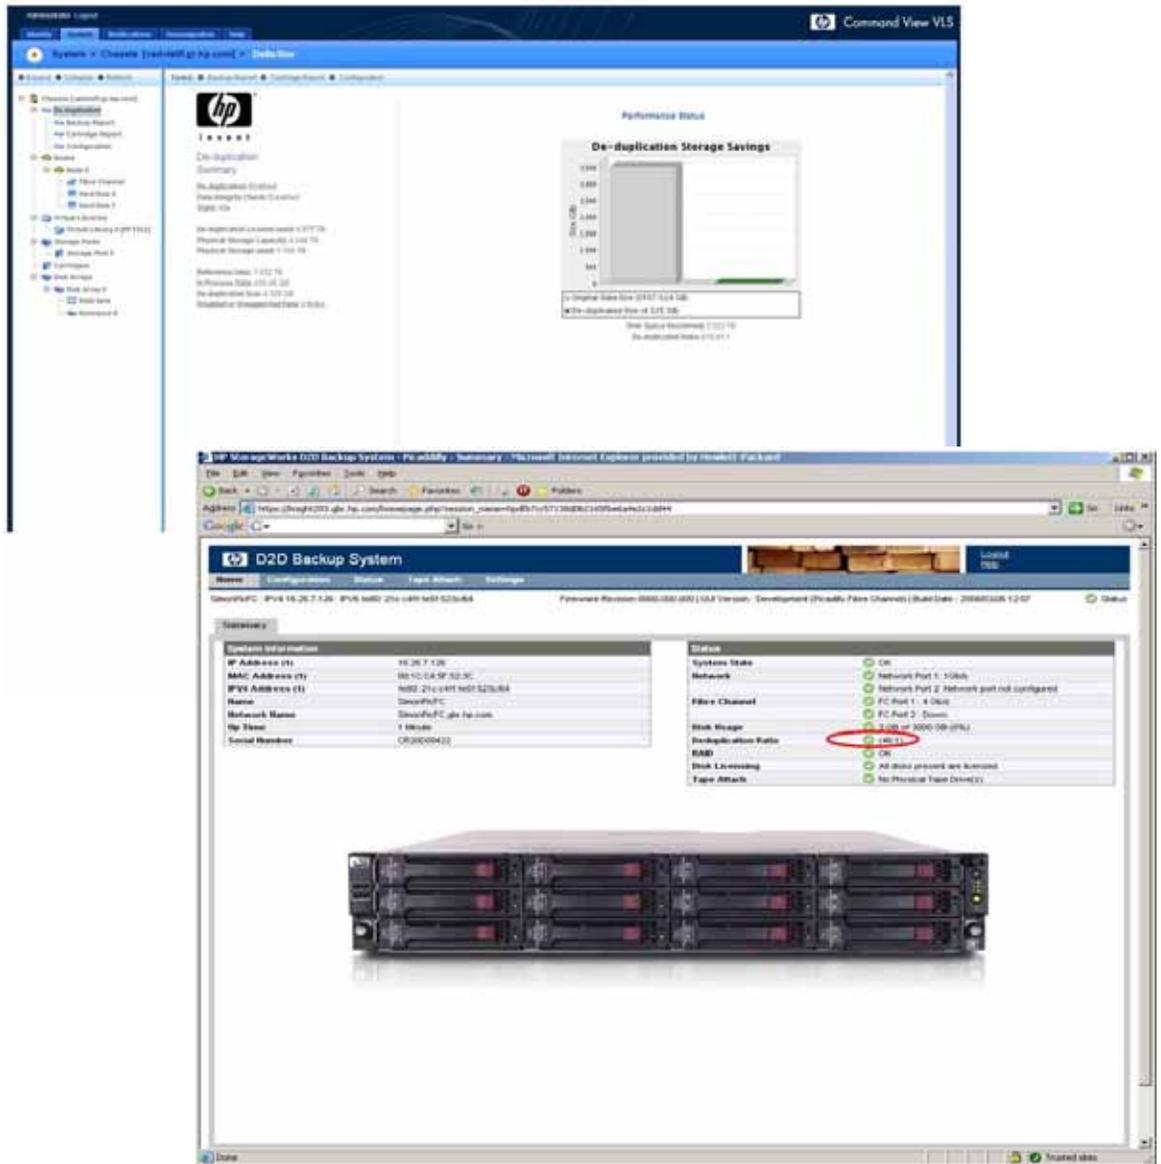
This whitepaper explains how HP deduplication technologies work in practice, the pros and cons of each approach, when to choose a particular type, and the type of low-bandwidth replication models HP plans to support.

Why HP for Deduplication?

- The HP Virtual Library System (VLS) incorporate Accelerated deduplication technology that delivers high-performance deduplication for enterprise customers. HP is one of the few vendors to date with an object level differencing architecture that combines the virtual tape library and the deduplication engine in the same appliance. Our competitors with object level differencing use a separate deduplication engine and VTL, which tends to be inefficient, as data is shunted between the two appliances, as well as expensive.
- HP D2D (Disk to Disk) Backup Systems use Dynamic deduplication technology that provides a significant price advantage over our competitors. The combination of HP patents allows optimal RAM and disk usage, intelligent chunking, and minimal paging. Together with the cost benefits of using HP industry-standard ProLiant servers sets a new price point for deduplication appliances.

HP D2D Backup Systems and VLS virtual libraries provide deduplication ratio monitoring as can be seen in the following screenshots.

Figure 1. Deduplication ratio screens on HP VLS and D2D devices



Introduction

Over recent years, virtual tape libraries have become the backbone of a modern data protection strategy because they offer:

- Disk-based backup at a reasonable cost
- Improved backup performance in a SAN environment because new resources (virtual tape drives) are easier to provision.
- Faster single file restores than physical tape
- Seamless integration into an existing backup strategy, making it low risk
- The ability to offload or migrate the data to physical tape for off-site disaster recovery or for long-term archiving

Because virtual tape libraries are disk-based backup devices with a virtual file system and the backup process itself tends to have a great deal of repetitive data, virtual tape libraries lend themselves particularly well to data deduplication. In storage technology, deduplication essentially refers to the elimination of redundant data. In the deduplication process, duplicate data is deleted, leaving only one copy of the data to be stored. However, indexing of all data is still retained should that data ever be required. Deduplication is able to reduce the required storage capacity since only the unique data is stored.

The amount of duplicate data that can be typically removed from a particular data type is estimated to be as follows:-

PACS	5%
Web and Microsoft office Data	30%
Engineering Data Directories	35%
Software code archive	45%
Technical Publications	52%
Database Backup	70% or higher

In the above example PACs are “Picture Archiving and Communication systems,” a type of data used in X-rays and medical imaging. These have very little duplicate data. At the other end of the spectrum, databases contain a lot of redundant data—their structure means that there will be many records with empty fields or the same data in the same fields.

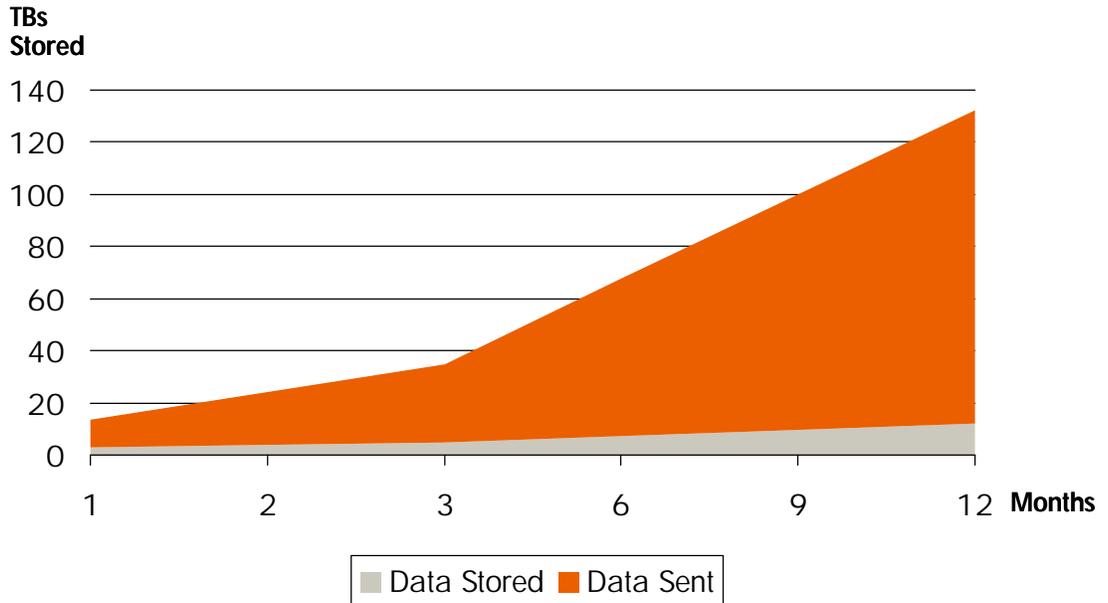
With a virtual tape library that has deduplication, the net effect is that, over time, a given amount of disk storage capacity can hold more data than is actually sent to it. To work deduplication needs a random access capability offered by disk based backup. This is not to say physical tape is dead, indeed tape is still required for archiving and disaster recovery, both disk and tape have their own unique attributes in a comprehensive data protection solution.

The capacity optimization offered by deduplication is dependent on:

- Backup policy, full, incremental
- Retention periods
- Data rate change

Figure 2. A visual explanation of deduplication

Why use Deduplication?



A word of caution

Some people view deduplication with the approach of "That is great! I can now buy less storage," but it does not work like that. Deduplication is a cumulative process that can take several months to yield impressive deduplication ratios. Initially, the amount of storage you buy has to be sized to reflect your existing backup tape rotation strategy and expected data change rate within your environment.

HP has developed deduplication sizing tools to assist with deciding the amount of storage capacity with deduplication that is required. However, these tools do rely on the customer having a degree of knowledge of how much the data change rate is in their systems.

[HP Backup Sizer Tool](#)

Deduplication has become popular because as data growth soars, the cost of storing data also increases, especially backup data on disk. Deduplication reduces the cost of storing multiple backups on disk. Deduplication is the latest in a series of technologies that offer space saving to a greater or lesser degree. To compare Deduplication with other data reduction, or space saving technologies please look at Appendix B.

A worked example of deduplication is illustrated as follows:

Figure 3. A worked example of deduplication for file system data over time

Example—1 TB file server backup

Retention policy

- 1 week, daily incrementals (5)
- 6 months, weekly fulls (25)

Data parameters

- Data compression rate = 2:1
- Daily change rate = 1%

(10% of data in 10% of files)

	Data sent from backup host	Data stored with deduplication
1 st daily full backup	1000 GB	500 GB
1 st daily incremental backup	100 GB	5 GB
2 nd daily incremental backup	100 GB	5 GB
3 rd daily incremental backup	100 GB	5 GB
4 th daily incremental backup	100 GB	5 GB
5 th daily incremental backup	100 GB	5 GB
1 st weekly full backup	1000 GB	25 GB
2 nd weekly full backup	1000 GB	25 GB
3 rd weekly full backup	1000 GB	25 GB
4 th weekly full backup	1000 GB	25 GB
5 th weekly full backup	1000 GB	25 GB
Total	25,500 GB	1,125 GB

~23:1 reduction in data stored

The example uses a system containing 1TB of backup data which equates to 500 GB of storage with 2:1 data compression for the first normal full backup. If 10% of the files change between backups, then a normal incremental backup would send about 10% of the size of the full backup or about 100 GB to the backup device. However, only 10% of the data actually changed in those files which equates to a 1% change in the data at a block or byte level. This means only 10 GB of block level changes or 5 GB of data stored with deduplication and 2:1 compression. Over time, the effect multiplies. When the next full backup is stored, it will not be 500 GB, the deduplicated equivalent is only 25 GB because the only block-level data changes over the week have been five times 5 GB incremental backups. Over the course of 6 months, a deduplication-enabled backup system uses the same amount of storage as less than a week of a normal backup system. Over a 6 month period, would provide a 23:1 effective savings in storage capacity. It also provides the ability to restore from further back in time without having to go to physical tape for the data. The key thing to remember here is that the deduplication ratio depends primarily on two things:

- What % of the data is changing between backups (% of data in % of files)
- How long is the retention period of the backups stored on disk

For example, a 0.5% daily change in the data in 10% of the files would yield a 50:1 deduplication ratio over one year of daily full backups. Obviously the % daily change rate is quite difficult to predict for complex systems, especially for applications like Exchange, SQL and Oracle so benchmarking is strongly advised.

Customer Benefits of Data Deduplication

What data deduplication offers to customers is:

- The ability to store dramatically more data **online (by online we mean disk based)**
- An increase in the range of Recovery Point Objectives (RPOs) available—data can be recovered from further back in time from the backup to better meet Service Level Agreements (SLAs). Disk recovery of a single files is always faster than tape
- A reduction of investment in physical tape by restricting its use more to a deep archiving and Disaster recovery usage model
- Deduplication can automate the disaster recovery process by providing the ability to perform site to site replication at a lower cost. Because 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 and is one of the most attractive propositions that deduplication offers. Customers who do not use disk based replication across sites today will embrace low-bandwidth replication, as it enables better disaster tolerance without the need and operational costs associated with transporting data off-site on physical tape. Replication is performed at a tape cartridge level

Figure 4. Remote site data protection BEFORE low-bandwidth replication

Remote site data protection before low bandwidth replication

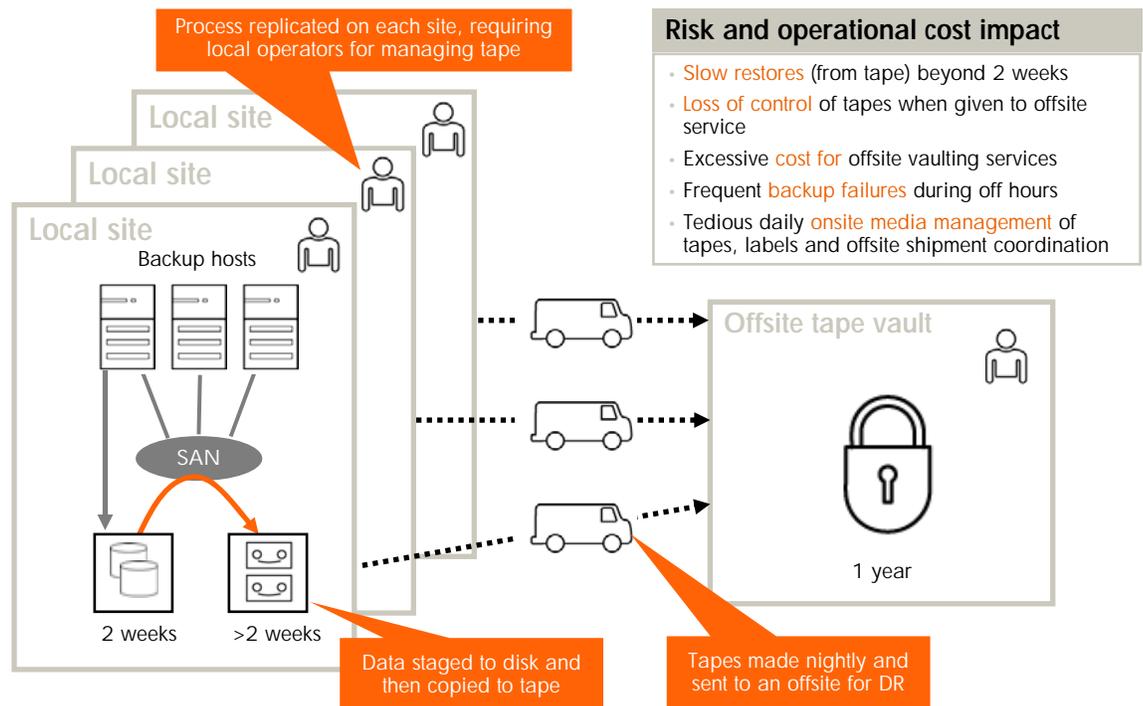
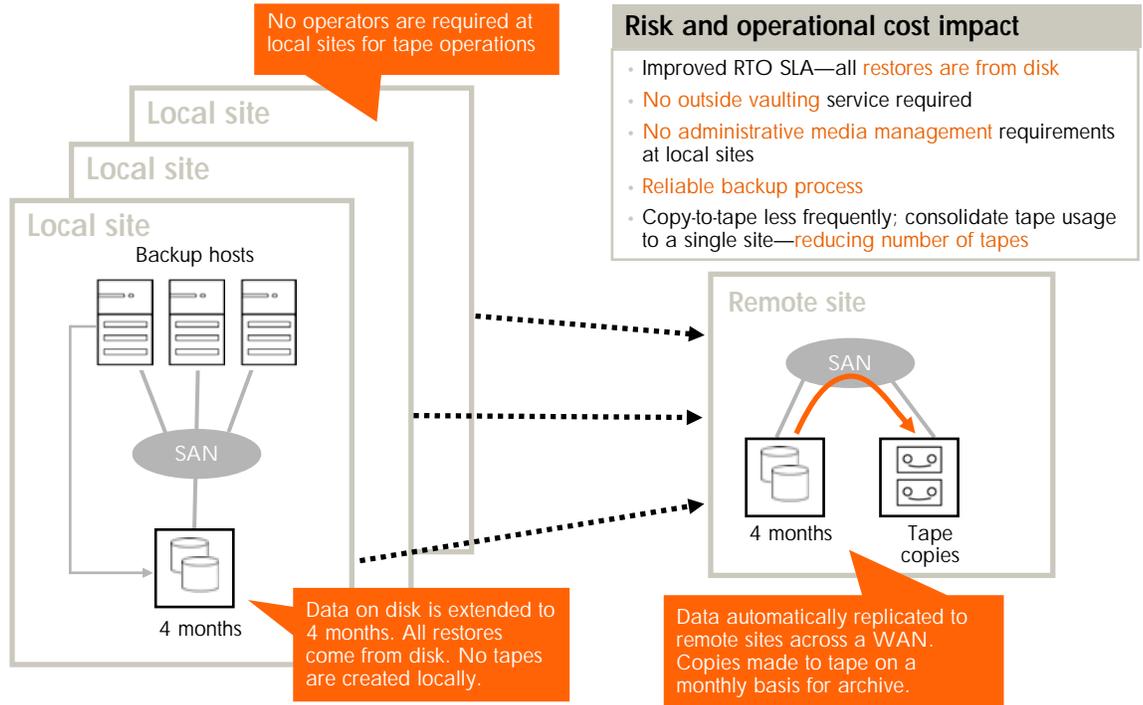


Figure 5. Remote site data protection AFTER low-bandwidth replication

Remote site data protection after low bandwidth replication



To show how much of an impact deduplication can have on replication times, take a look at the following Figure 6. This model also takes into account a certain overhead in control information that has to be sent site to site as well as the data deltas themselves. Currently without deduplication the full amount of data has to be transferred between sites and in general this requires high bandwidth links such as GbE or Fibre Channel. With Deduplication only the delta changes are transferred between sites and this reduction allows lower bandwidth links such as T3 or OC12 to be used at lower cost. The following example illustrates the estimated replication times for varying amounts of change. Most customers would be happy with a replication time of say 2 hours between sites using say a T3 link. The feed from HP D2D backup systems or HP Virtual Library systems to the replication link is one or more GbE pipes.

Figure 6. Replication times with and without deduplication

Estimated Time to Replicate Data for a 1TB Backup Environment @ 2:1

		Link Type		
	Data Sent	T1	T3	OC12
Link Rate (66% efficient)		1.5 Mb/s	44.7 Mb/s	622.1 Mb/s
Backup Type	Without dedupe			
Incremental	50 GB	4.5 days	3.8 hrs	16 min
Full	500 GB	45.4 days	1.6 days	2.7 hrs
Change Rate	With dedupe			
0.5%	13.1 GB	29 hrs	59 min	4.3 min
1.0%	16.3 GB	35 hrs	73 min	5.3 min
2.0%	22.5 GB	49 hrs	102 min	7.3 min

A word of caution

An initial synchronization of the backup device at the primary site and the one at the secondary site must be performed. Because the volume of data that requires synchronizing at this stage is high, a low-bandwidth link will not suffice. Synchronization can be achieved in three different ways:

- Provision the two devices on the same site and use a feature such as local replication over high-bandwidth fibre channel links to synchronize the data. Then ship one of the libraries to the remote site
- Install the two separate devices at separate sites, perform initial backup at Site A. Copy the backup from Site A to physical tape, then transfer the physical tapes to site B and import them. When the systems at both sites are synchronized, start low-bandwidth replication between the two
- After initial backup at site A allow a multi-day window for initial synchronization allowing the two devices to copy the initial backup data over a low-bandwidth link

Understanding Customer Needs for Data Deduplication

Both large and small organizations have remarkably similar concerns when it comes to data protection. What differs is the priority of their issues.

Figure 7. Common challenges with data protection amongst remote offices, SMEs and large customers

Common challenges

Environment

Needs

Data center



- Handle explosive data growth
- Meet and maintain backup windows
- Achieve greater backup reliability
- Accelerate restore from “tape” (inc virtual tape)
- Manage remote site data protection

Remote office



- Overcome a lack of dedicated IT resources
- Manage data growth
- Maintain backup application, file and OS independence
- Spend less time managing backups

SME

Different priorities are what have led HP to develop two distinct approaches to data deduplication. For example:

- Large enterprises have issues meeting backup windows, so any deduplication technology that could slow down the backup process is of no use to them. Medium and Small enterprises are concerned about backup windows as well but to a lesser degree
- Most large enterprise customers have Service Level Agreements (SLAs) pertaining to restore times—any deduplication technology that slows down restore times is not welcome either
- Many large customers back up hundreds of terabytes per night, and their backup solution with deduplication needs to scale up to these capacities without degrading performance. Fragmenting the approach by having to use several smaller deduplication stores would also make the whole backup process harder to manage
- Conversely remote offices and smaller organizations generally need an easy approach—a dedicated appliance that is self-contained—at a reasonable cost
- Remote offices and SMEs do not want or need a system that is infinitely scalable, and the cost associated with linearly scalable capacity and performance. They need a single engine approach that can work transparently in any of their environments

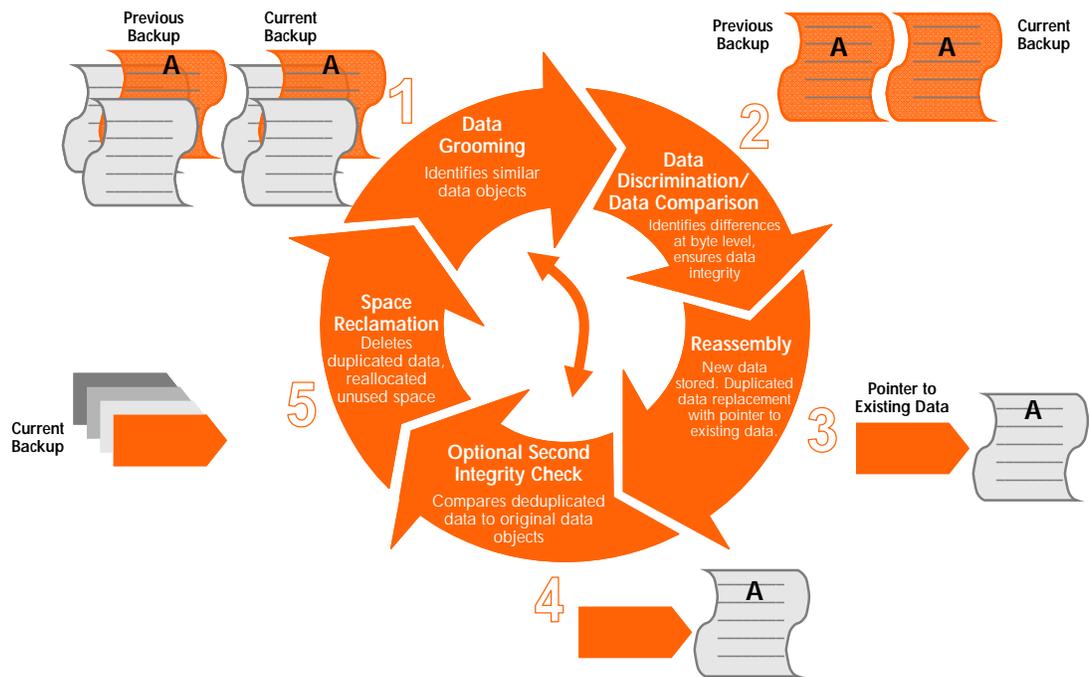
HP Accelerated Deduplication for the Large Enterprise Customer

HP Accelerated deduplication technology is designed for large enterprise data centers. It is the technology HP has chosen for the HP StorageWorks Virtual Library Systems. Accelerated deduplication has the following features and benefits:

- Utilizes object-level differencing technology with a design centered on performance and scalability
- Delivers fastest possible backup performance—it leverages post-processing technology to process data deduplication as backup jobs complete, deduplicating previous backups whilst other backups are still completing.
- Delivers fastest restore from recently backed up data—it maintains a complete copy of the most recent backup data and eliminates duplicate data in previous backups
- Scalable deduplication performance—it uses distributed architecture where performance can be increased by adding additional nodes
- Flexible replication options to protect your investment

Figure 8. Object-level differencing compares only current and previous backups from the same hosts and eliminates duplicate data by means of pointers. The latest backup is always held intact.

HP Accelerated Deduplication



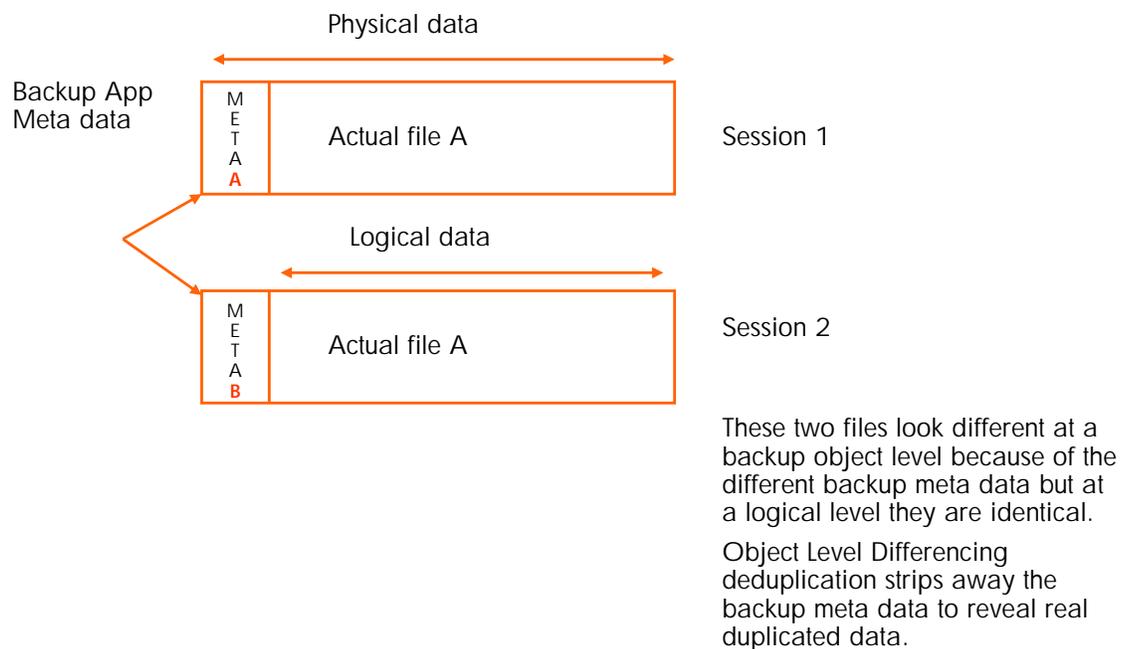
How Accelerated Deduplication Works

When the backup runs the data stream is processed as it is stored to disk assembling a content database on the fly by interrogating the meta data attached by the backup application. This process has minimal performance impact.

1. After the first backup job completes, tasks are scheduled to begin the deduplication processing. The content database is used to identify subsequent backups from the same data sources. This is essential, since the way object-level differencing works is to compare the current backup from a host to the previous backup from that same host.

Figure 9. Identifying duplicated data by stripping away the meta data associated with backup formats, files and databases

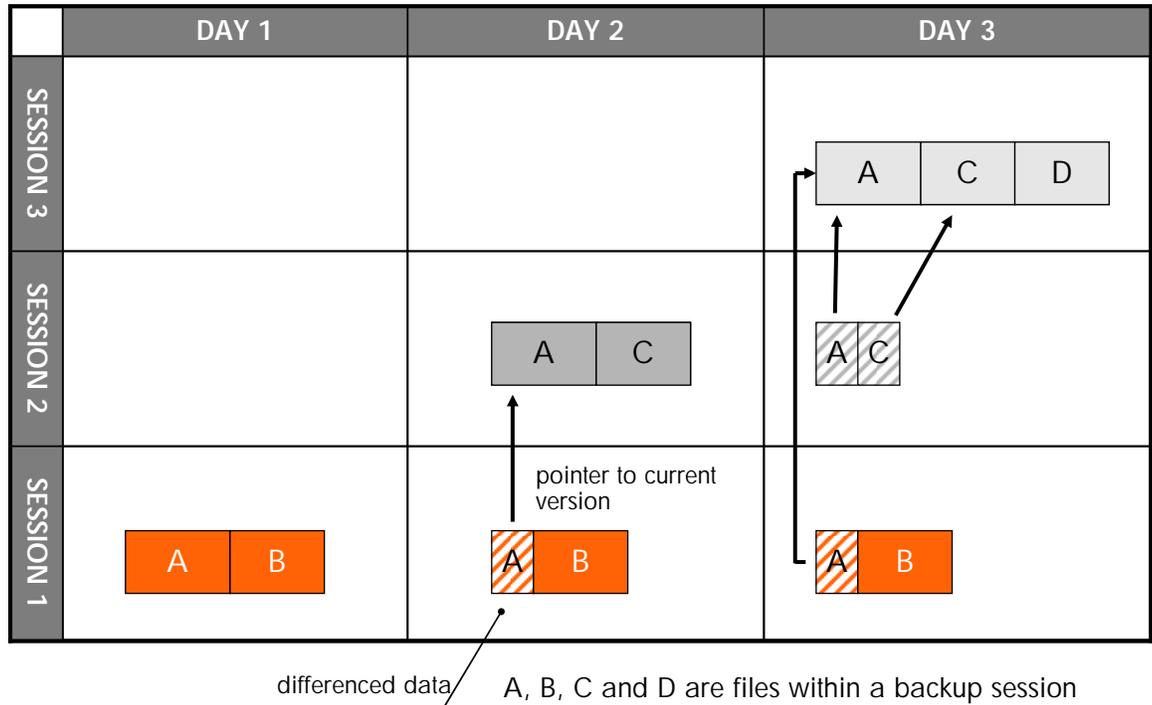
Object Level Differencing strips away the meta data to reveal real duplication



2. A data comparison is performed between the current backup and the previous backup from the same host. There are different levels of comparison. For example, some backup sessions are compared at an entire session level. Here, data is compared byte-for-byte between the two versions and common streams of data are identified. Other backup sessions compare versions of files within the backup sessions. Note that within Accelerated deduplication's object-level differencing, the comparison is done AFTER the backup meta data and file system meta data has been stripped away. (See the example in the following Figure 10) This makes the deduplication process much more efficient but relies on an intimate knowledge of both the backup application meta data types and the data type meta data (file system file, database file, and so on).
3. When duplicate data is found in the comparison process, the duplicate data streams in the oldest backup are replaced by a set of pointers to a more recent copy of the same data. This ensures that the latest backup is always fully contiguous, and a restore from the latest backup will always take place at maximum speed.

Figure 10. With object level differencing the last backup is always fully intact. Duplicated objects in previous backups are replaced with pointers plus byte level differences.

HP Accelerated Data Deduplication – Details



In the preceding diagram, backup session 1 had files A and B. When backup session 2 completed and was compared with backup session 1, file A was found and a byte-level difference calculated for the older version. So in the older backup (session1), file A was replaced by pointers plus difference deltas to the file A data in backup session 2. Subsequently, when backup session 3 completes, it is compared with backup session 2 and file C is found to be duplicated. Hence a difference and a pointer is placed in backup session 2 pointing to the file C data in backup session 3, also at the same time the original pointer to file A in Session1 is readjusted to point to the new location of file A. This is to prevent multiple hops for pointers when restoring older data. So the process continues, every time comparing the current backup with the previous backup. Each time a difference plus pointer is written, storage capacity is saved. This process allows the deduplication to track even a byte level change between files.

4. Secondary Integrity Check—before a backup tape is replaced by a deduplicated version with pointers to a more recent occurrence of that data, a byte-for-byte comparison can take place comparing the original backup with the “reconstructed” backup, including pointers to ensure that the two are identical. Only when the compare succeeds will the original backup tape be replaced by a version including pointers. This step is optional. See Figure 9 Step 4.
5. Space reclamation occurs when all the free space created by replacing duplicate data with pointers to a single instance of the data is complete. This can take some time and results in used capacity being returned to a free pool on the device.

Replication can take place from Step 3 because the changed data is available to be replicated even before the space has been reclaimed.

HP Accelerated deduplication:

- Will scale up to hundreds of TB
- Has no impact on backup performance, since the comparison is done after the backup job completes (post process)
- Allows more deduplication “compute nodes” to be added to increase deduplication performance and ensure the post processing is complete before the backup cycle starts again.
- Yields high deduplication ratios because it strips away meta data to reveal true duplication, and does not rely on data chunking.
- Provides fast bulk data restore and tape cloning for recently backed up data—maintains the complete most recent copy of the backup data but eliminates duplicate data in previous backups.

Issues Associated with Object-Level Differencing

The major issue with object-level differencing is that the device has to be knowledgeable in terms of backup formats and data types to understand the Meta data. HP Accelerated deduplication will support a subset of backup applications and data types at launch.

Additionally, object-level differencing compares only backups from the same host against each other, so there is no deduplication across hosts, but the amount of common data across different hosts can be quite low.

What Makes HP Accelerated Deduplication unique?

The object-level differencing in HP Accelerated deduplication is unique in the marketplace. Unlike hash-based techniques that are an all-or-nothing method of deduplication, object-level differencing applies intelligence to the process, giving users the ability to decide what data types are deduplicated and allowing flexibility to reduce the deduplication load if it is not yielding the expected or desired results. HP Object-level differencing technology is also the only deduplication technology that can scale to hundreds of terabytes with no impact on backup performance, because the architecture does not depend on managing ever increasing amounts of Index tables, as is the case with Hash based chunking. It is also well suited for larger scaleable system since it is able to distribute the deduplication workload across all the available processing resources and can even have dedicated nodes purely for deduplication activities.

- HP Accelerated deduplication will be supported on a range of backup applications:
 - HP Data Protector
 - Symantec NetBackup
 - Tivoli Storage Manager
 - Legato Networker
- HP Accelerated deduplication will support a wide range of file types:
 - Windows 2003
 - Windows Vista
 - HP-UX 11.x
 - Solaris standard file backups
 - Linux Redhat
 - Linux SuSe
 - AIX file backups
 - Tru64 file backups

- HP Accelerated deduplication will support database backups over time:
 - Oracle RMAN
 - Hot SQL Backups
 - Online Exchange
 - MAPI mailbox backups

For the latest more details on what Backup software and data types are supported with HP Accelerated Deduplication please look at the HP Enterprise Backup Solutions compatibility guide at <http://www.hp.com/go/ebs>

HP Accelerated deduplication technology is available by license on HP StorageWorks Virtual Library Systems (models 6000, 9000, and 12000). The license fee is per TB of user storage (before compression or deduplication takes effect).

Figure 11. Pros and cons of HP Accelerated Deduplication

Pros & Cons of HP Accelerated Deduplication

PRO	CON
<ul style="list-style-type: none"> • Does not restrict backup rate – since data is processed after the backup has completed. • Faster restore rate – “forward referencing pointers” allow rapid access to data. • Can handle datasets > 100TB without having to partition backups – no hashing table dependencies. • Can selectively compare data likely to match, increasing performance further – higher deduplication ratios. • Best suited to large Enterprise VTLs. 	<ul style="list-style-type: none"> • Has to be ISV format aware and data type aware, content coverage will grow over time. • May need additional compute nodes to speed up post processing deduplication in scenarios with long backup windows. • Needs to cache 2 backups in order to perform post process comparison. So additional disk capacity equal to the size of the largest backup needs to be sized into the solution.

At ingest time when the tape content database is generated there is a small performance overhead (< 0.5%) and there is a small amount of disk space required to hold this database (much less than the hash tables in the hash based chunking deduplication technology). Even if this content database were completely destroyed it would still be possible to maintain access to the data because the pointers are still fully intact and held within the re-written tape formats.

HP object level differencing also has ability to provide selective deduplication by content type, and in the future could be used to index content providing content addressable archive searches.

The question often arises “What happens if deduplication is not complete by the time the same backup from the same host arrives?” Typically the deduplication process takes about 2 x as long as the backup process for a given backup, so as long as a single backup job does not take > 8 hours

this will not occur. In addition the multi-node architecture ensures that each node is load balanced to provide 33% of its processing capabilities to deduplication whilst still maintaining the necessary performance for backup and restore. Finally additional dedicated 100% deduplication compute nodes can be added if necessary.

Let us now analyze HP's second type of deduplication technology—Dynamic deduplication, which uses hash-based chunking.

HP Dynamic Deduplication for Small and Medium IT Environments

HP Dynamic deduplication is designed for customers with smaller IT environments. Its main features and benefits include:

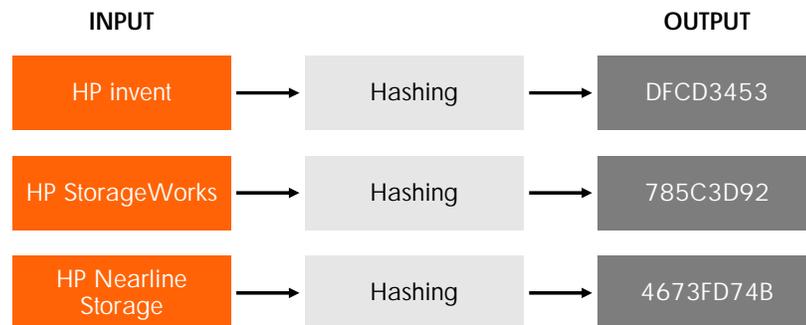
- Hash-based chunking technology with a design center around compatibility and cost
- Low cost and a small RAM footprint
- Independence from backup applications
- Systems with built-in data deduplication
- Flexible replication options for increased investment protection.

Hash-based chunking techniques for data reduction have been around for years. Hashing consists of applying an algorithm to a specific chunk of data and yielding a unique fingerprint of that data. The backup stream is simply broken down into a series of chunks. For example, a 4K chunk in a data stream can be “hashed” so that it is uniquely represented by a 20-byte hash code. See Figure 13

Figure 12. Hashing technology

Hashing Technology

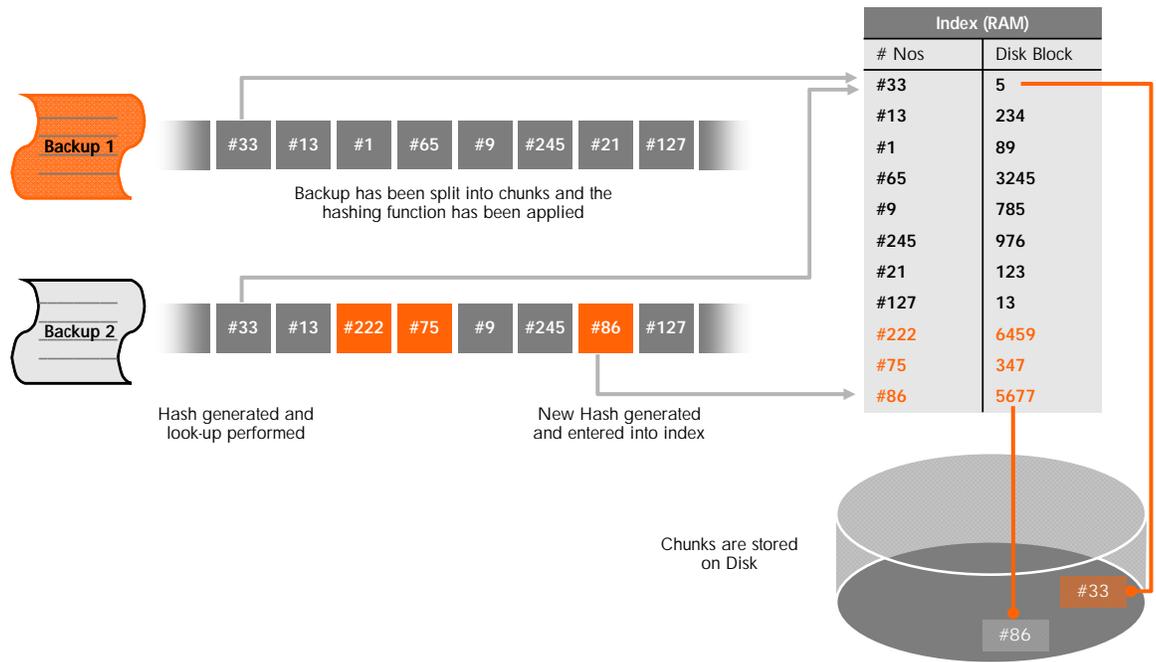
- **“in-line”** = deduplication on the fly as data is ingested using “hashing” techniques
- **“hashing”** = is a reproducible method of turning some kind of data into a (relatively) small number that may serve as a digital “fingerprint” of the data



The larger the chunks, the less chance there is of finding an identical chunk that generates the same hash code—thus, the deduplication ratio will not be as high. The smaller the chunk size, the more

efficient the data deduplication process, but then a larger number of indexes are created, which leads to problems storing enormous numbers of indexes (see the following example and Glossary).

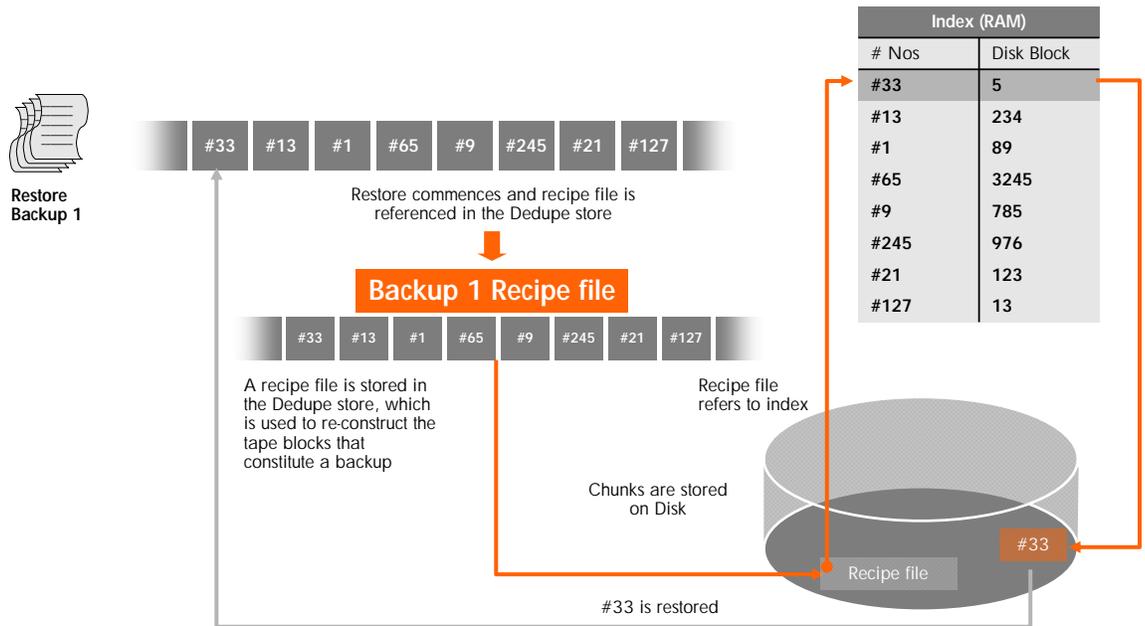
Figure 13. How hash based chunking works



How Dynamic Deduplication Works

1. As the backup data stream enters the target device (in this case the HP D2D2500 or D2D4000 Backup System), it is chunked into nominal 4K chunks against which the SHA-1 hashing algorithm is run. These results are placed in an "index" (hash value) and stored in RAM in the target D2D device. The hash value is also stored as an entry in a "recipe" file which represents the backup stream, and points to the data in the deduplication store where the original 4K chunk is stored. This happens in real time as the backup is taking place. Step 1 continues for the whole backup data stream.
2. When another 4K chunk generates the same hash index as a previous chunk, no index is added to the index list and the data is not written to the deduplication store. An entry with the hash value is simply added to the "recipe file" for that backup stream pointing to the previously stored data, so space is saved. Now as you scale this up over many backups there are many instances of the same hash value being generated, but the actual data is only stored once, so the space savings increase.
3. Now let us consider backup 2 in Figure 13. As the data stream is run through the hashing algorithm again, much of the data will generate the same hash index codes as in backup 1—hence, there is no need to add indexes to the table or use storage in the deduplication store. In this backup, some of the data has changed. In some cases (#222, #75, and #86), the data is unique and generates new indexes for the index store and new data entries into the deduplication store.
4. And so the hashing process continues ad infinitum until as backups are overwritten by the tape rotation strategy certain hash indexes are no longer required, and so in a housekeeping operation they are removed.

Figure 14. How hash-based chunking performs restores



5. On receiving a restore command from the backup system, the D2D device selects the correct recipe file and starts sequentially re-assembling the file to restore.
 - a. Read recipe file.
 - b. Look up hash in index to get disk pointer.
 - c. Get original chunk from disk.
 - d. Return data to restore stream.
 - e. Repeat for every hash entry in the recipe file.

Issues Associated with Hash-Based Chunking

The main issue with hash-based chunking technology is the growth of indexes and the limited amount of RAM storage required to store them. Let us take a simple example: if we have a 1TB backup data stream using 4K chunks, and every 4K chunk produces a unique hash value. This equates to 250 million 20-byte hash values or 5GB of storage.

If we performed no other optimization (for example, paging of indexes onto and off disk), then the appliance would need 5GB of RAM for every TB of deduplicated unique data. Most server systems cannot support much more than 16GB of RAM. For this reason, hash-based chunking cannot easily scale to hundreds of terabytes.

Most lower-end to mid-range deduplication technologies use variations on hash-based chunking, but with additional techniques to reduce the size of the indexes generated, reducing the amount of RAM required, but generally at the expense of some deduplication efficiency or performance. If the index management is not efficient, it will slow the backup down to unacceptable levels or miss many instances of duplicate data. The other option is to use larger chunk sizes to reduce the size of the index. As mentioned earlier, the downside of this is that deduplication will be less efficient. These algorithms can also be adversely affected by non-repeating data patterns that occur in some backup software tape formats. This becomes a bigger issue with larger chunk sizes.

HP has developed a unique innovated technology leveraging work from HP Labs that dramatically reduces the amount of memory required for managing the index without sacrificing performance or deduplication efficiency. Not only does this technology enable low-cost high performance disk backup systems, but it also allows the use of much smaller chunk sizes to provide more effective data deduplication which is more robust to variations in backup stream formats or data types.

Restore times can be slow with hash-based chunking. As you can see from figure 14, to recover a 4K piece of data from a hash-based deduplication store requires a reconstruction process. The restore can take longer than it did to back up.

Finally you may here the term “hashing collisions”—this means that 2 different chunks of data produce the same hash value, which obviously undermines the data integrity. The chances of this happening are remote to say the least. HP Labs calculated

Using a TWENTY BYTE (160 bit) hash such as SHA1, the time required for a hashing collision to occur is 100,000,000,000,000 years, based on the backing up 1TB of data per working day.

Even so, HP Dynamic deduplication adds a further Cyclic Redundancy Checksum (CRC) at a tape record level that would catch the highly unlikely event of a hash collision.

Despite the above limitations, deduplication using hash-based chunking is a well-proven technology and serves remote offices and medium sized businesses very well. The biggest benefit of hash-based chunking is that it is totally data format-independent and it does not have to be engineered to work with specific backup applications and data types. The products using the hash based deduplication technology still have to be tested with the various backup applications but the design approach is generic.

HP is deploying Dynamic deduplication technology on its latest D2D Backup Systems, which are designed for remote offices and small to medium organizations.

HP D2D 2500 and 4000 Backup Systems come with deduplication as standard with no additional licensing costs.

Figure 15. Pros and cons of hash-based chunking deduplication

Pros & Cons of HP Dynamic Deduplication

PRO	CON
<ul style="list-style-type: none">• Deduplication performed at backup time• Can instantly handle any data format• Significant processing overhead, but keeping pace with processor developments.• Fast search, algorithms already proven to aid hash detection• Low storage overhead – don't have to hold complete backups (TBs) for post analysis• Best suited to smaller size VTLs	<ul style="list-style-type: none">• Can restrict ingest rate (backup rate) if not done efficiently and could slow backups down.• Restore time may be longer than object level differencing deduplication because of data regeneration process.• Concerns over scalability when using very large hash indexes. For data sets > 100TB may have to start "partitioning" backups to ensure better hash index management.

What makes HP Dynamic Deduplication technology unique are algorithms developed with HP Labs that dramatically reduce the amount of memory required for managing the index, and without sacrificing performance or deduplication effectiveness. Specifically, this technology:

- Uses far less memory by implementing algorithms that determine which are the most optimal indexes to hold in RAM for a given backup data stream
- Allows the use of much smaller chunk sizes to provide more effective data deduplication which is more robust to variations in backup stream formats or data types
- Provides intelligent storage of chunks and recipe files to limit disk I/O and paging
- Works well in a broad range of environments since it is independent of backup software format and data types

Low-Bandwidth Replication Usage Models

The second main benefit of deduplication is the ability to replicate the changes in data on site A to a remote site B at a fraction of the cost because high-bandwidth links are no longer required. A general guideline is that a T1 link is about 10% of the cost of a 4Gb FC link over the same distance. Low-bandwidth replication will be available on both D2D and VLS products. Upto 2 GbE ports will be available for replication on D2D devices and 1 GbE port per node will be available on the VLS products.

HP will support three topologies for low bandwidth replication:

- Box --> Box
- Active <-> Active
- Many ---> One

The unit of replication is a cartridge. On VLS, it will be possible to partition slots in a virtual library replication target device to be associated with specific source replication cartridges.

Figure 16. Active <-> Active replication on HP VLS and D2D systems with deduplication

Accelerated Deduplication Replication Example Use Case – Active/Active

- Generally datacenter-to-datacenter replication, with each device performing local backups and also acting as the replication store for the other datacenter

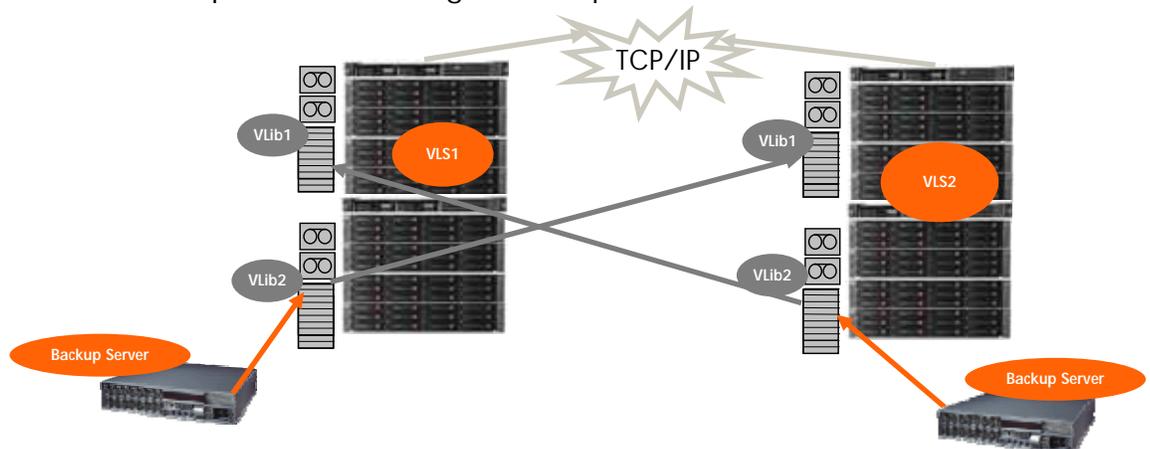
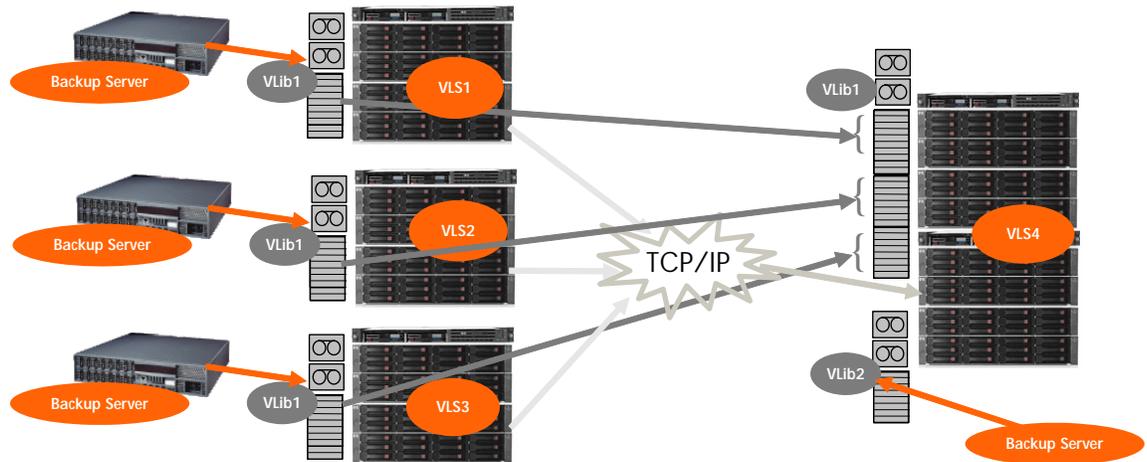


Figure 17. Many-to-one replication on HP VLS and D2D systems with deduplication

Accelerated Deduplication Replication Example Use Case – Many to One

- Can divide up a single destination target into multiple slots ranges to allow many-to-one without needing a separate replication library for each one



Initially it will not be possible for D2D devices to replicate into the much larger VLS devices, since their deduplication technologies are so different, but HP plans to be able to offer this feature in the near future.

What will be possible is to replicate multiple HP D2D250 into a central D2D4000 or replicate smaller VLS6200 models into a central VLS 12000 (See Figure 18)

Deduplication technology is leading us to the point where many remote sites can replicate data back to a central data center at a reasonable cost, removing the need for tedious off-site vaulting of tapes and fully automating the process—saving even more costs.

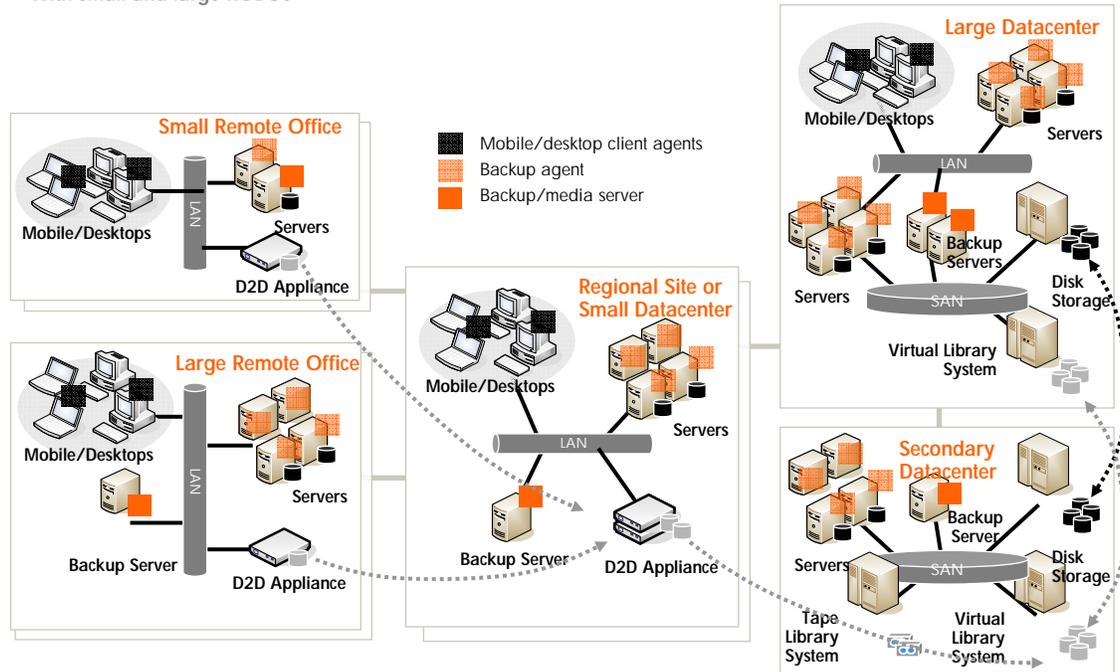
This ensures

- The most cost effective solution is deployed at each specific site
- The costs and issues associated with off site vaulting of physical tape are removed
- The whole Disaster recovery process is automated
- The solution is scalable at all sites

Figure 18. Enterprise Deployment with replication across remote and branch offices back to data centers

Enterprise Deployment

With small and large ROBOs



Why HP for Deduplication?

Deduplication is a powerful technology and there are many different ways to implement it, but most vendors offer only one method and, as we have seen, no one method is best in all circumstances. HP offers a choice of deduplication technologies depending on your needs. HP does not pretend that "one size fits all."

Choose HP Dynamic deduplication for small and mid-size IT environments because it offers the best technology footprint for deduplication at a price point that is affordable. Flexible replication options further enhance the solution.

Choose HP Accelerated deduplication for Enterprise data centers where scalability and backup performance are paramount. Flexible replication options further enhance the solution.

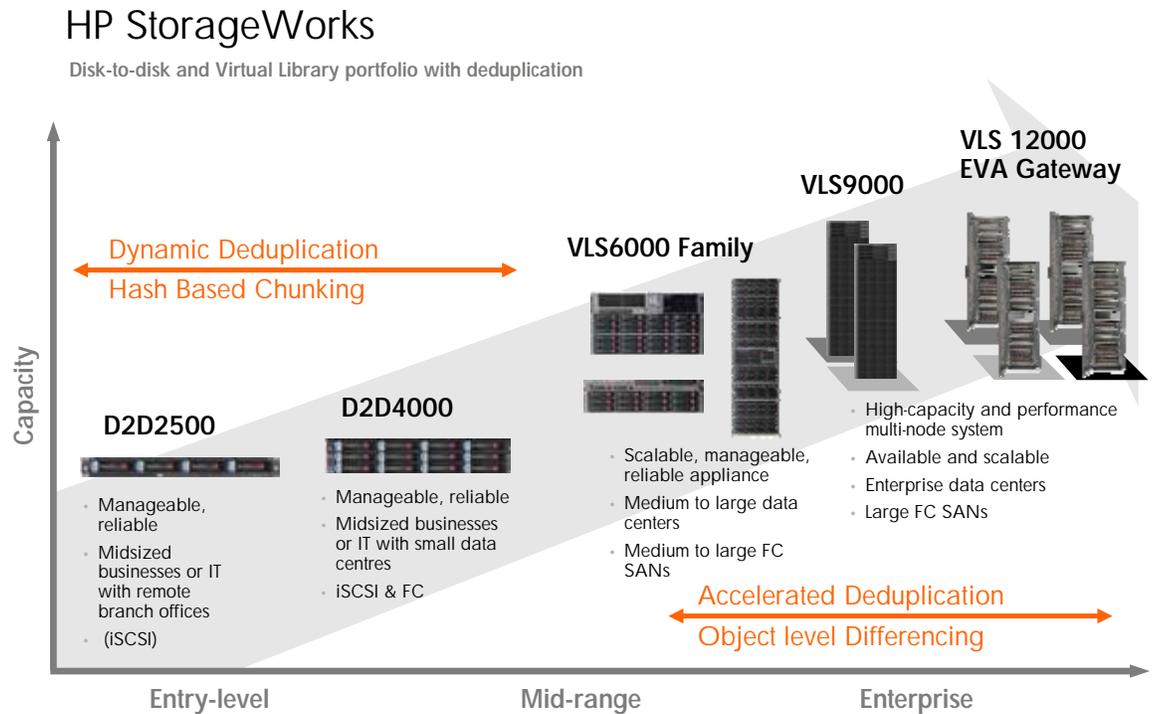
The scalability issues associated with hash-based chunking are addressed by some competitors by creating multiple separate deduplication stores behind a single management interface, but what this does is create "islands of deduplication," so the customer sees reduced benefits and excessive costs because the solution is not inherently scalable.

At the data center level, the major competitors of HP using object-level differencing have used "bolt-on" deduplication engines with existing virtual tape library architectures and have not integrated the deduplication engine within the VTL itself. This leads to data being moved back and forth between the virtual library and the deduplication engine, which is very inefficient.

Deduplication Technologies Aligned with HP Virtual Library Products

HP has a range of disk-based backup products with deduplication starting with the entry-level D2D2500 at 2.25TB user unit for small businesses and remote offices, right up to the VLS12000 EVA Gateway with capacities over 1 PB for the high-end enterprise data center customer. They emulate a range of HP Physical tape autoloaders and libraries.

Figure 19. HP disk-based backup portfolio with deduplication



The HP StorageWorks D2D2500 and D2D4000 Backup Systems support HP dynamic deduplication. These range in size from 2.25TB to 7.5TB and are aimed at remote offices or small enterprise customers. The D2D2500 has an iSCSI interface to reduce the cost of implementation at remote offices, while the D2D4000 offers a choice of iSCSI or 4Gb FC.

The HP StorageWorks Virtual Library Systems are all 4Gb SAN-attach devices which range in native user capacity from 4.4TB to over a petabyte with the VLS9000 and VLS12000 EVA Gateway. Hardware compression is available on the VLS6000, 9000 and 12000 models, achieving even higher capacities. The VLS9000 and VLS12000 use a multi-node architecture that allows the performance to scale in a linear fashion. With eight nodes, these devices can sustain a throughput of up to 4800MB/sec at 2:1 data compression, providing the SAN hosts can supply data at this rate. HP Virtual Library Systems will deploy the HP Accelerated deduplication technology.

Summary

Data deduplication technology represents one of the most significant storage enhancements in recent years, promising to reshape future data protection and disaster recovery solutions. Deduplication offers the ability to store more on a given amount of storage and enables replication using low-bandwidth links, both of which improve cost effectiveness.

HP offers two complementary deduplication technologies for different customer needs:

- Accelerated deduplication (with object-level differencing) for high-end enterprise customers who require:
 - Fastest possible backup performance
 - Fastest restore
 - Most scalable solution in terms of performance and capacity
 - Multi-node low bandwidth replication
 - Highest deduplication ratios
 - Wide range of replication models
- Dynamic deduplication (with hash-based chunking) for mid size organizations and remote offices that require:
 - Lower cost and a smaller footprint
 - An integrated deduplication appliance with lights-out operation
 - Backup application and data type independent for maximum flexibility
 - Wide range of replication models

This whitepaper explained how deduplication technologies of HP work in practice, the pros and cons of each approach, when to choose a particular type and the type of low-bandwidth replication models HP plans to support.

The HP Virtual Library System (VLS) incorporate Accelerated deduplication technology that scales for large multi-node systems and delivers high-performance deduplication for enterprise customers.

HP D2D (Disk to Disk) Backup Systems use Dynamic deduplication technology that provides a significant price advantage over our competitors. The combination of HP patents allowing optimal RAM usage (RAM footprint) with minimal new hash values being generated on similar backup streams. HP D2D backup systems with integrated deduplication set a new price point for deduplication devices.

Appendix A—Glossary of Terminology

Source-based Deduplication

Where data is deduplicated in the host(s) prior to transmission over the storage network. This generally tends to be a proprietary approach.

Target-based Deduplication

This is where the data is deduplicated in a Target device such as a virtual tape library and is available to all hosts using that target device.

Hashing

This is a reproducible method of turning some kind of data into a (relatively) small number that may serve as a digital "fingerprint" of the data.

Chunks

This is a method of breaking down a data stream into segments (chunks), and on each chunk the hashing algorithm is run.

SHA-1

Secure hashing algorithm 1. For example SHA-1 can enable a 4K chunk of data to be uniquely represented by a 20-byte hash value.

Object-Level Differencing

Is a general IT description that describes a process that has an intimate knowledge of the data that it is handling—down to logical format level. Object-level differencing deduplication means the deduplication process has an intimate knowledge of the backup application format, the file types being backed up (for example, Windows file system, exchange files, and SQL files). This intimate knowledge allows file comparisons at a byte level to remove duplicated data.

Box-to-Box

Replication from a Source to Destination in one direction.

Active-Active

Replication from a Source device on Site A to a Target Device on Site B and vice versa.

Many-to-One

Replication from multiple sources to a single destination device.

Deduplication ratio

The reduction in storage required for a backup (after several other backups have taken place). Figures between 10:1 and 300:1 have been quoted by different vendors. The ratio is highly dependent on:

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

Space Reclamation

With all deduplication devices time is required to free up space that was used by the duplicated data and return it to a "free pool" Because this can be quite time consuming it tends to occur in off peak periods.

Post Processing

This is where the deduplication is done AFTER the backup completes to ensure there is no way the deduplication process can slow down the backup and increase the backup window required.

In-Line

This is where the deduplication process takes place REAL TIME as the backup is actually taking place. Depending on implementations this may or may not slow the backup process down.

Multi-thread

Within HP Object Level differencing the compare and space reclamation processes are run with multiple paths simultaneously to ensure faster execution times.

Multi-node

HP VLS9000 and VLS12000 products scale to offer very high performance levels—up to eight nodes can run in parallel, giving throughput capabilities up to 4800MB/sec at 2:1 compression ratio. This multi-node architecture is fundamental to HPs Accelerated deduplication technology because it allows maximum processing power to be applied to the deduplication process.

Appendix B—Deduplication compared to other data reduction technologies

Technology description	Pro	Con	Comments
Deduplication —Advanced technique for efficiently storing data by referencing existing blocks of data that have been previously stored, and only storing new data that is unique.	Two fold benefits Space savings of between 10 and 100:1 being quoted Further benefit of low bandwidth replication	Can slow backup down if not implemented efficiently. Hash based technologies may not scale to 100s of TB Object Level differencing technologies need to be multi format aware which takes time to engineer	Deduplication is by far the most impressive disk storage reduction technology to emerge over recent years. Implementation varies by vendors. Benchmarking highly recommended
Single Instancing —Is really deduplication at a file level	Available as part of the Microsoft file system and as a feature of the file system of a Netapp filer. System based approach to space savings	Will not eliminate redundancy within a file, only if two files are exactly the same For example adding files to a PST file, or adding a slide to a presentation.	Limited use
Array-based 'snapshots' capture changed blocks on a disk LUN	Used primarily for fast roll-back to a consistent state using "image recovery"—not really focused on storage efficiency.	Does not eliminate redundant data for the changed blocks Captures any change made by the file system—example does not distinguish between real data and deleted/free space on disk	Well established. Generally used for quick recovery to a known point in time
Incremental Forever backups —recreate a full restore image from just one full backup and lots of incrementals	Minimizes the need for frequent full backups and hence allows for smaller backup windows	More focused at time savings than really at space savings	Generally only works with file system backups not database based backups
Compression —software or hardware	Fast (if done in hardware), slower if done in software. Well established and understood	Maximum space savings are generally 2:1	Can be used in addition to deduplication

For more information

www.hp.com/go/tape

www.hp.com/go/D2D

www.hp.com/go/VLS

www.hp.com/go/deduplication

[HP StorageWorks customer success stories](#)

© Copyright 2008 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.

Linux is a U.S. registered trademark of Linus Torvalds. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. UNIX is a registered trademark of The Open Group.

4AA1-9796ENW, May 2008

