

## About HiberniaEvros

Hibernia Evros is an IT solutions and services company specialising in the design, supply, implementation and support of mission critical ICT infrastructure.

Offering the highest levels of service and technical expertise, backed by over 20 years of experience, Hibernia Evros keeps over 800 Irish organisations operating effectively.

By working closely with an extensive network of partners, our Professional Services group are able to design the best solution available to meet the needs of each particular customer.

Once implemented, our Managed Services group provide a range of technical support services to ensure your technology continues to support your business.

## Working with Us

We have built our unrivalled reputation on offering the highest standards of service available in Ireland today. We have achieved this by following three simple rules in everything we do:

### Being Technically Excellent

We strive to offer the best technologies implemented and supported by the best technical practitioners working in Ireland today.

### Putting the Customer First

We take time to understand your business, so we can deliver a solution designed to suit you – not us!

### Being easy to deal with

We won't waste your time with unnecessary red tape. Instead we offer a friendly, no nonsense service.

Data Domain deduplication storage systems dramatically reduce the amount of disk storage needed to retain and protect enterprise data. By identifying redundant files and data as they are being stored, Data Domain systems provide a storage footprint that is 10x-30x smaller, on average, than the original data set. Backup data can then be efficiently replicated and retrieved over existing networks for streamlined disaster recovery and consolidated tape operations.

Data Domain is a strategic infrastructure partner of Hibernia Evros and we have unrivalled experience in Ireland of designing and implementing Data Domain technologies to allow our customers to enjoy the benefits of this leading data deduplication technology.

## Why deduplicate data?

Eliminating redundant data can significantly shrink storage requirements and improve bandwidth efficiency. Because the cost of primary storage is continually reducing, enterprises typically store many versions of the same information so that new work can re-use old work. But this can result in significant additional costs in storage – both operational and backup.

Deduplication lowers storage costs since fewer disks are needed, and shortens backup/recovery times since there can be far less data to transfer. In the context of backups there is a great deal of duplicate data. The same data keeps getting stored over and over again consuming a lot of unnecessary storage space (disk or tape), electricity) and bandwidth (for replication), creating a chain of cost and resource inefficiencies within the organisation

## How does data deduplication work?

Deduplication segments the incoming data stream, uniquely identifies the data segments, and then compares the segments to previously stored data. If an incoming data segment is a duplicate of what has already been stored, the segment is not stored again, but a reference is created to it. If the segment is unique, it is stored on disk.

For example, a file or volume that is backed up every week creates a significant amount of duplicate data. Deduplication algorithms analyse the data and can store only the compressed, unique change elements of that file. This process can provide an average of 10-30 times or greater reduction in storage capacity requirements, with average backup retention policies on normal enterprise data. This means that companies can store 10TB to 30TB of backup data on 1 TB of physical disk capacity, which has huge economic benefits.

## Is deduplication easy to implement?

Data Domain has made it very easy to implement by creating a fast, application-independent storage system (attachable as a file server over Ethernet or a VTL over Fibre Channel). No client software or other configuration is required. As a result, Data Domain deduplication is invisible to backup and recovery and other nearline storage processes. It works easily with various data movers and workloads, including non-backup data like e-mail archives, reference data and engineering revision libraries. More flexibility means more consolidation is possible using less physical infrastructure.

## What data deduplication rates are expected?

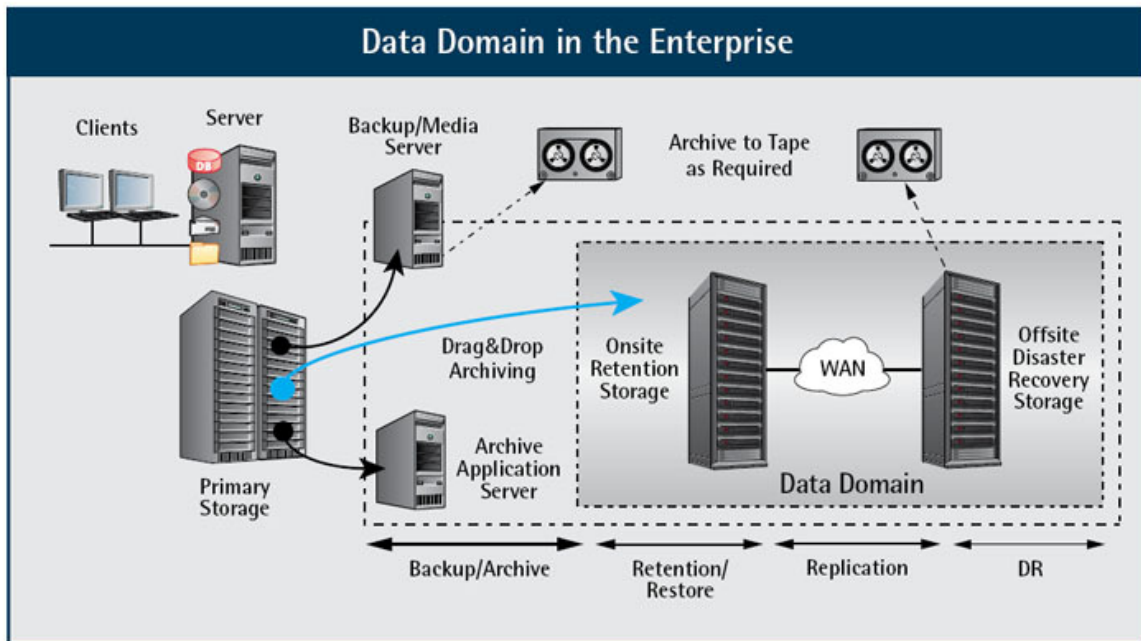
Redundancy will vary by application, frequency of version capture and retention policy. Significant variables include the rate of data change (few changes mean more data to deduplicate), the frequency of backups (more fulls makes compression effect higher), the retention period (longer retention means more data to compare against), and the size of the data set (more data, more to deduplicate).

Generally, Data Domain enables 2x-4x data reduction on an initial full backup, 6x-7x reduction on subsequent file-level incrementals, and 50x-60x reduction on subsequent full backups

**Leading the way in inline deduplication**

Inline deduplication means the data is deduplicated before it is written to disk (inline). Post-process deduplication analyses and reduces data after it has been stored to disk. Data Domain has class leading inline deduplication technology.

Inline deduplication is the most efficient and economic method of deduplication as it significantly reduces the raw disk capacity needed in the system since the full, not-yet-deduplicated data set is never written to disk. If replication is supported as part of the inline deduplication process, the time-to-DR (disaster recovery) is reduced far beyond all other methods. This is possible as the system does not need to wait to absorb the entire data set and then deduplicate it before it can begin replicating to the remote site.



Data Domain systems deduplicate data inline at speeds of up to 1.4 TB per hour (22.4 TB per hour for the DDX Array). Data is stored on-site for immediate restores and longer term retention on disk. The deduplicated data can also be replicated over the WAN to a remote site for disaster recovery purposes, eliminating the need for tape backups, or for consolidating tape backups to a central location. Data Domain deduplication storage systems provide a full range of replication options, including selective replication, many-to-one and bi-directional replication.

**Improve off-site replication and DR with Data Deduplication**

The impact that deduplication has on replication and disaster recovery windows can be profound. Deduplication results in significantly less data being transmitted to keep the DR site up to date. In turn, this means that less expensive WAN links are required.

Within a deduplicated environment the replication process is also significantly speeded up as there is less data to send. The inline deduplication approach made possible by the use of Data Domain technology will yield the most aggressive and efficient results as it happens during the backup, significantly improving the time by which there is a complete restore point at the DR site, or improving the time to DR readiness.

Typically less than 1% of a full backup is actually new, unique deduplicated data sequences that can be sent over a WAN immediately upon the start of the backup. Aggressive cross-site deduplication, when multiple sites replicate to the same destination, can add additional value by deduplicating across all backup replication streams and all local backups. Unique deduplicated segments previously transferred by any remote site, or held in local backup, are then used in the deduplication process to further improve network efficiency by reducing the data to be vaulted. In other words, if the destination system already has a data sequence that came from a remote site or a local backup and that same sequence is created at another remote site, it will be identified as redundant by the Data Domain system before it consumes bandwidth traveling across the network to the destination system. All of the data collected at the destination site can be safely moved off-site to a single location or multiple DR sites.

To see how Data Domain can make a real difference in your environment, visit [www.dedupecalculator.com](http://www.dedupecalculator.com)