



## FAQs

Virtual Tape Library (VTL) with Deduplication

# Frequently Asked Questions

Virtual Tape Library (VTL) with Deduplication enables fast, reliable restores, saving considerable time and resources.



## **WHAT IS DYNAMIC SOLUTIONS INTERNATIONAL (DSI) VIRTUAL TAPE LIBRARY (VTL) WITH DEDUPLICATION?**

The DSI Virtual Tape Library (VTL) with Deduplication is the industry's market-leading virtual tape and LAN-based deduplication solution, unmatched in performance and scalability. With virtual tape and NAS interfaces (NFS/SMB), backups complete faster and more reliably, with little or no change needed to the backup environment. Sophisticated physical tape integration and global data deduplication complete the solution. The DSI Virtual Tape Library (VTL) increases the speed and reliability of standard third-party backup application environments by leveraging disk to emulate industry-standard tape libraries. Integrated global deduplication minimizes the amount of storage needed for backup operations, providing significant cost savings.

## **WHY DO I NEED THE DSI VIRTUAL TAPE LIBRARY (VTL) WITH DEDUPLICATION?**

Due to the explosive growth of data from databases, email applications, and other business applications, backup administrators are finding it increasingly difficult, if not impossible, to complete backups within allotted backup windows. Furthermore, media and device failures are commonplace, requiring valuable time and resources to keep the backup environment up-and-running.

With virtual tape and NAS interfaces (NFS/SMB), backups complete faster and more reliably, with little or no change needed to the backup environment.

Virtual Tape Library (VTL) with Deduplication enables fast, reliable restores and eliminates the need to re-run backup jobs due to tape errors, saving considerable time and resources. Because this solution uses disk to back up data, it eliminates the media and mechanical errors that can occur with physical tapes and drives. In addition, Virtual Tape Library (VTL) can emulate more tape drives than your physical tape library so you can run more backup streams simultaneously. This enables you to more easily complete backups within shrinking backup windows.

The DSI Virtual Tape Library (VTL) with Deduplication can be installed at remote sites, consolidating backup data at a central location utilizing an existing tape infrastructure. You can easily add new virtual tape drives and libraries, or increase backup jobs and backup streams, without the expense and hassle of adding new hardware.

## HOW DOES THE DEDUPLICATION ENGINE WORK?

The deduplication engine works within the DSI Virtual Tape Library (VTL) to eliminate redundant data without impacting established backup windows, while minimizing replication time and storage requirements. The deduplication engine scans backup data at the block level, analyzes the data, and determines whether data is unique or has already been stored in the deduplication repository. Only single instances of data are stored. The original virtual tape is replaced with a virtual index tape (VIT) pointing to previously stored data, freeing space for more data. The DSI VTL with Deduplication offers a choice of deduplication options: Inline, postprocess, turbo deduplication, or no deduplication.

## WHEN WOULD I NEED THESE DEDUPLICATION OPTIONS?

Engineered with users in mind, the solution provides the flexibility to select the desired deduplication method at the tape level, based on unique and changing requirements:

- Inline deduplication, introduced in the DSI370 Series of this solution, minimizes storage requirements, reducing them by as much as 40%. It is ideal for small storage configurations or environments where immediate replication is desired. For the highest performance levels, a recommended best practice is to use the product's flexible deduplication policies to leverage post-process deduplication for the initial backup, and then switch to inline deduplication for subsequent backups.
- Post-process deduplication is ideal when you deploy external EVD nodes. As its name implies, it occurs after the backup process completes, thus it can be scheduled to run at any time.
- Concurrent deduplication (a.k.a. Turbo) runs in a similar fashion to post-processing, but it starts as soon as the first virtual tape has been written, running concurrently with backups. Deduplication engines start working immediately, making full use of available CPU resources, which is highly suitable for clustered environments. Because replication starts sooner, data can be quickly recovered from a remote site. Turbo deduplication does not read the virtual tapes again after backup is done while post-processing deduplication does. It can reduce the storage load as a result.
- No deduplication can be used on data that does not deduplicate effectively or data that is exported to physical tape. Examples include image files and data that is pre-compressed or encrypted. By selectively "turning off" deduplication, you can save cycles and focus on deduplicating the data that will yield the highest overall ratios and value.

## I AM CURRENTLY USING THE POST-PROCESS/TURBO DEDUPLICATION FUNCTIONALITY OF THIS SOLUTION. CAN I USE INLINE DEDUPLICATION AS WELL?

Yes. Inline deduplication is compatible with existing deduplication processes. Inline deduplication can coexist with existing post-process and turbo processes, wherein three processes can run at the same time on different tapes. Tapes using turbo deduplication can be moved seamlessly to inline deduplication as needed, and vice-versa. This unique flexibility allows you to adjust to your changing needs without any impact on backup processes.

## HOW DOES DATA DEDUPLICATION BENEFIT REPLICATION?

Without deduplication, all of a virtual tape's new backup data is sent over the wire, resulting in multiple terabytes of duplicate data driving up bandwidth requirements and costs. Since much of the backup data is redundant, only new/unique information needs to be replicated. By identifying unique data, the deduplication engine minimizes the bandwidth needed to replicate backup data, minimizing costs and making replication practical for most organizations.

## WHAT IS THE VALUE OF GLOBAL DEDUPLICATION?

Since local data deduplication makes replication more affordable, you can easily and cost-effectively consolidate backup data from multiple remote offices to a central location such as a data center. However, there may still be duplicate data between remote offices. Global deduplication ensures that redundant data from remote offices is not replicated to the central repository, thus minimizing space requirements.

**WHAT IS THE POSITION OF THE DSI VTL WITH DEDUPLICATION IN THE MARKETPLACE?** The DSI Virtual Tape Library (VTL) is the leader in the virtual tape interface segment of the deduplication appliance market, with recognition from industry press and analysts. The enterprise-level scalability, flexibility, and features

eclipse other vendors, offering the ability to adapt to customer requirements. A wide range of industry-leading strategic business partners attest to the quality of the solution. Moreover, customer acceptance is the truest measure of product leadership, and more customers choose the DSI Virtual Tape Library (VTL) than any other solution of its kind.

### **DO I NEED TO CHANGE MY CURRENT BACKUP SOFTWARE?**

No. The DSI Virtual Tape Library (VTL) enhances the performance of your current backup application, so you can decrease your current backup window without having to modify job configurations. The solution works with the most common backup and restore applications available in the market.

### **HOW DOES VTL WITH DEDUPLICATION SUPPORT LONG-TERM ARCHIVING?**

To meet retention requirements, data on virtual tape or NAS share can be exported directly to physical tape for long-term archiving (D2D2T). Data can continue to be copied to physical tape using the backup application's "copy" function. The unique open architecture of this solution enables integration with local and remote physical tape resources, for the best possible return on investment (ROI).

### **WHAT KIND OF TAPE LIBRARIES AND TAPE DRIVES DOES THIS SOLUTION EMULATE?**

The DSI Virtual Tape Library (VTL) can emulate more tape libraries and tape drives than any other vendor. For each node, you can emulate up to 128 virtual tape libraries, 1,024 virtual drives, and 1,000,000 (1 million) virtual tapes.

A complete list of supported vendors can be found on the DSI Certification Matrix at <https://www.dynamicsolutionsinternational.com>

### **WHAT KIND OF LAN-BASED FILE INTERFACE DOES OPTIMIZED BACKUP & DEDUPLICATION SUPPORT?**

The DSI Virtual Tape Library (VTL) offers SMB or NFS network share over IP for NetBackup and Backup Exec backup applications. This easy-to-use file interface is compatible with all source data, seamlessly integrating with major backup software, database utilities, archiving applications, virtual machine data, and manual file copy methods. There is little or no need to change D2D backup applications or file and data archiving processes. Backup performance and reliability are drastically improved, and secondary storage capacity requirements are minimized.

### **HOW DO I MOVE DATA FROM VIRTUAL TAPE OR NAS SHARE TO PHYSICAL TAPE?**

You can use the DSI Virtual Tape Library (VTL) to copy data to physical tape, or you can use any backup software that offers copy-to-tape functionality. Using backup software enables you to leverage the same processes as you would with physical tape. However, this puts an additional load on your backup server and the traffic between it and the virtual tape library. The DSI Virtual Tape Library (VTL) with Deduplication enables you to keep barcodes in sync locally and remotely, letting you proactively manage the tape export process.

### **CAN I MAKE MULTIPLE COPIES OF TAPES SIMULTANEOUSLY?**

Yes. You can export up to five virtual tapes to physical tapes at a time.

### **DO I NEED TO RESERVE A LARGE QUANTITY OF STORAGE FOR MY BACKUPS?**

No. The DSI Virtual Tape Library (VTL) includes a "Capacity on Demand" feature that allows you to create and allocate space as needed. This can save considerable disk space without affecting system performance. Inline deduplication allows you to back up directly to the repository further reducing the capacity requirement.

### **WHAT IS THE AUTOMATED TAPE CACHING OPTION?**

With Automated Tape Caching, tapes always appear to be inside the DSI solution and are visible to the backup software regardless of whether the data resides on disk or tape. This gives the backup application direct access to data at any time. The DSI Virtual Tape Library (VTL) acts as a transparent cache to the physical tape library. Based on extremely flexible migration policies (age of data, time of day, disk space, end of backup, etc.), the data is staged to the virtual tape cache and copied to physical tape after the backup completes.

Automated Tape Caching also provides very flexible space reclamation policies (immediately upon migration, after a specified retention period, when you run out of space, etc.). When this is integrated with the deduplication process, you can trigger



If you need to recover files, you can use the import function to directly access the physical tape for immediate recovery.

deduplication and replication instead of deletion from disk. This provides greater restore availability and performance.

### **CAN I IMPORT PHYSICAL TAPES TO THIS SOLUTION?**

Yes. A key advantage of the DSI Virtual Tape Library (VTL) is that you can protect data on your existing physical tapes by importing them into your virtual tape system. If you need to recover files, you can use the import function to directly access the physical tape for immediate recovery. The import function allows you to:

- Copy contents of a physical tape onto a virtual tape
- Directly access data from physical tape without copying the entire tape
- Recycle a physical tape
- Import data from virtual tapes being used by another appliance

### **CAN I USE MY CURRENT BACKUP SOFTWARE TO EXPORT TO PHYSICAL TAPE?**

Yes. The DSI Virtual Tape Library (VTL) allows you to export data from virtual to physical tape using any NDMP v.4 supported backup software, without sending the data through the backup server. With this option enabled, a DSI appliance acts as an NDMP server, centralizing management by eliminating locally attached tape devices from each NAS device. When backup occurs, data is moved from the NAS device directly to the appliance.

### **WHAT HAPPENS IF AN APPLIANCE RUNS OUT OF CAPACITY?**

When a DSI appliance reaches full capacity, all running backup jobs will be given 'end-of-tape' or 'no space left on device' indicators, and backup jobs will fail due to a lack of media space. You have three options:

1. You can move data to physical tapes using the backup software and reuse the old virtual tapes.
2. You can move data to physical tapes using this solution and create new tapes in the freed-up space.
3. You can add storage capacity while the DSI solution is running to create more virtual tapes or expand the NAS share.

### **MY CURRENT BACKUP SOFTWARE TOUTS BACKUP-TO-DISK CAPABILITIES. WHY SHOULD I CONSIDER THE DSI VTL WITH DEDUPLICATION?**

Backup-to-disk performance is usually impacted by an intervening file system and progressive fragmentation. The DSI Virtual Tape Library (VTL) with Deduplication utilizes raw storage streaming to achieve higher performance levels on more cost-efficient storage. Global enterprise-class deduplication decreases the storage requirements and enables more efficient replication.

Long-term archive requirements are most cost-effectively addressed

by physical tape. The DSI VTL offers many physical tape integration features that others lack, including a Secure Tape option that uses the Advanced Encryption Standard (AES) 256-bit algorithm. With this option you can create one or more keys that can be used to encrypt the data when it is exported to physical tape, and decrypt it when it is imported back to virtual tape. The data on the tape remains fully encrypted.

### **IS THIS SOLUTION COMPLIANT WITH THE STANDARDS OF THE FEDERAL INFORMATION PROCESSING STANDARDS?**

Yes. The DSI VTL supports compliance with the US Government Federal Information Processing Standards (FIPS).

### **DOES THIS SOLUTION SUPPORT CLUSTER CONFIGURATIONS?**

Yes. For high availability (HA), VTL with Deduplication can scale up to four pairs (eight nodes). If a machine fails, another will automatically take over its workload. The DSI EVD node supports N+1 clustering of up to 4 active nodes, with a total repository size up to 2PB.

One, two, or four active nodes can support a single deduplication repository for enterprise-level performance and capacity scalability. When paired with a standby node, the standby system (+1) will automatically take its place in the event of failure to ensure continuity. This allows you to start with the smallest possible configuration that meets your exact needs, and expand incrementally as your data requirements grow.

### **WHICH CONNECTIVITY PROTOCOLS ARE SUPPORTED?**

VTL with Deduplication supports a variety of protocols, such as FC, iSCSI, NDMP, NFS, and SMB, for unparalleled performance. A single VTL node can back up over 39TB of data in an hour over four 16GB FC ports.

### **CAN I MANAGE MULTIPLE SERVERS FROM A SINGLE CONSOLE?**

Yes. A centralized console enables you to create a multi-node configuration that allows you to manage up to eight DSI VTL servers. This can include a combination of single servers and/or failover pairs, offering the following capabilities:

- Log in once: Single sign-on allows you to log into all of the servers in the group with a single user ID and password
- Deploy Deduplication Nodes
- Create/delete groups
- Add/remove members
- Run reports for all servers in a group. This includes standard reports that are generated on each server in the group and contain data specific to that server. You can also run a consolidated Group Disk Space Usage Report that includes every server in the group.
- Consolidated Event Log/Attention required monitoring: Event log displays events from all servers in chronological order.
- Monitor common properties, including compression, SNMP, storage monitoring triggers, and tape caching.

### **I WANT TO UPGRADE MY CURRENT TAPE DRIVE FORMAT TO EXPAND CAPACITY. CAN THIS SOLUTION HELP ME?**

Yes. A tape consolidation option allows you to export multiple virtual tapes to a single physical tape in a tape library. This maximizes tape usage by enabling the conversion of smaller-capacity virtual tape to higher-capacity physical tape (i.e. DLT to LTO). You can import all of your old tapes into the virtual tape library and consolidate the data onto a new larger tape. Tape consolidation can support up to 1,000 virtual tapes per physical tape.



## **WHAT ARE THE AVAILABLE FORM FACTORS OF THE DSI VIRTUAL TAPE LIBRARY (VTL) WITH DEDUPLICATION ?**

Designed with an organization's size and unique needs in mind, the solution is available in several options, including software and hardware:

- DSI Restore: For VMWare infrastructures.
- DSI VTL All-in-one appliances are ideal for IBM and Unisys environments. This solution integrates with IBMi & Unisys Clearpath systems and storage to deliver fast, reliable backup, restore, and remove DR for data assurance.
- DSI EVO & EVD Nodes for Midmarket and Enterprises. HA cluster deduplication gateways that integrate with certified SAN storage in midmarket and enterprise environments.

## **DOES THE DSI VIRTUAL TAPE LIBRARY (VTL) SUPPORT CASCADED REPLICATION?**

Yes. Cascaded replication is supported for VTL interface. You can configure cascaded replication for deduplication policies. Moreover, you can add exporting to physical tapes as one of the options after each replication, so you not only will have a virtual replica but also a physical copy at each replication target.

## **CAN I REPLICATE ONE TAPE TO TWO SITES AT THE SAME TIME?**

Yes. You can configure parallel replication in a deduplication policy.

## **HOW MANY WAYS CAN I DEPLOY VTL WITH DEDUPLICATION?**

The DSI Virtual Tape Library (VTL) with Deduplication can be deployed as a pre-built and preconfigured appliance, or as a preconfigured VMWare virtual appliance.

## **HOW CAN I SECURE MY BACKUP DATA FROM UNWANTED ACCESS?**

VTL with Deduplication offers extensive security options such as strong complex passwords, user password expiration, and user account lockout to protect organizational data from unwanted access. All data on a DSI VTL is encrypted, whether its in-flight, at-rest, or during physical tape exports.

Information in this document is provided "As Is" without warranty of any kind, and is subject to change without notice by Dynamic Solutions International (DSI), which assumes no responsibility for any errors or claims herein. Copyright 2018 Dynamic Solutions International. All rights reserved.