



Using Dynamic Solutions International's Virtual Tape Library (VTL) as a Backup Target for IBM i

ABSTRACT

For too long, the standard method of backing up IBM i (AS/400) systems has been with a locally attached tape drive or library. The result is reduced performance due to limited backup streams, excessive manual tape handling and expensive off-site data archiving. By introducing the Dynamic Solutions International (DSI) Virtual Tape Library (VTL) solution into an IBM i environment, organizations can non-disruptively replace tape backup in a disk-to-disk (D2D) format to accelerate backup/recovery processes and improve operational efficiency. The DSI VTL provides significant business and IT advantages for IBM i environments, including integration with native backup tools, IPL bare-metal recovery from virtual or physical tape and direct physical tape support with encryption to maximize data protection. This reduces operational and media costs while increasing performance levels to meet stringent application recovery service level agreements (SLAs) and to maximize return on investment (ROI).



INTRODUCTION

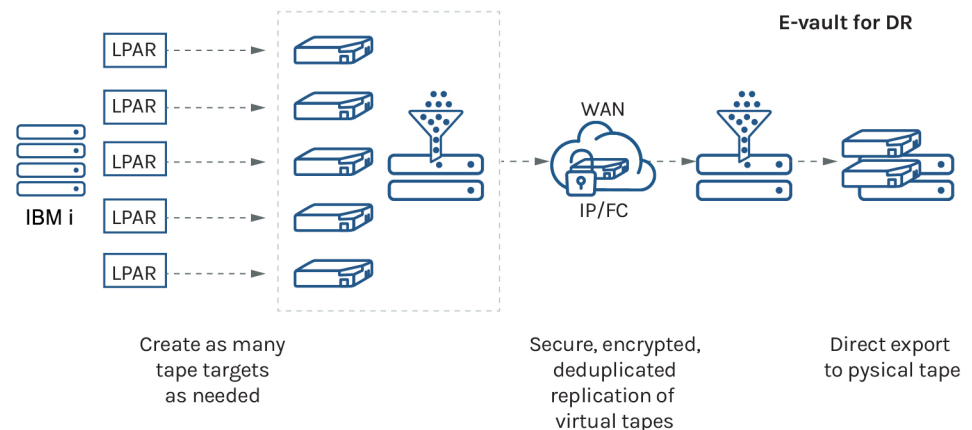
In today's data-driven business environment, customers face increased challenges in protecting their vital data from loss, theft, corruption, and disaster. Traditional backup operations constantly reproduce data for protection and recovery purposes; therefore the amount of data continues to increase and IT costs keep rising. Even though disk prices are lower each year and tape drive and SAN performance improves, coping with exponential data growth remains a significant challenge for most organizations.

Traditionally, the standard method of backing up IBM i systems has been with a locally attached tape drive or library. However, performance is inhibited due to limited backup streams, excessive manual tape handling and insufficient offsite data security. In addition, the IBM i backup environment is often an isolated IT island falling outside the rest of the backup operations.

With the introduction of the DSI Virtual Tape Library (VTL) solution, traditional tape backup is replaced by virtual tape in a disk-to-disk (D2D) format. The DSI VTL is a complete, integrated solution that improves the speed and performance of IBM i backup and recovery, enhances security via encryption to tape, and simplifies data movement. Implementing the DSI VTL enables high-performance restore, eliminates tape transport, and offers tape encryption and replication to meet secure offsite vaulting requirements.

Because the DSI VTL incorporates its own disk resources, it doesn't consume valuable IBM i storage, so there is no impact on internal IBM i disk consumption. Emulation of IBM Tape Libraries and Tape drives combined with native support for IBM i backup tools, makes it easy for backup administrators to manage their backup infrastructure. No additional knowledge is required. All operations and process are transparent and continue as normal. The only difference is that backup and restore is faster, and data is stored more reliably and efficiently.

NATIVE IBM I SUPPORT WITH FALCONSTOR VTL



Benefits include:

- Reduce total cost of ownership (TCO)
- Use virtual tape backups to start up IPL to boot IBM i
- Retain more data online with data deduplication
- Replicate data across sites to eliminate tape shipment
- Stream data quickly to physical tape as needed
- Encrypt data onto tape without the need for IBM i software or hardware-based library managed encryption (LME)



ABOUT IBM i

Companies rely on an efficient IT infrastructure to support business-critical applications. They need to know that their systems and business can meet the highest service levels defined by their business units and can be adapted to handle every new business opportunity. IBM i (formerly known as i5/OS) offers a highly scalable and virus-resistant architecture with a proven reputation for business resiliency.

Over the years IBM has transitioned from OS/400 to i5/OS to IBM i. When IBM introduced POWER5, OS/400 was renamed i5/OS. When the POWER6 platform became available in January 2008, IBM announced a major new release named i5/OS V6R1. Later that same year, the name was changed to IBM i to disassociate any presumed dependency on the POWER5 hardware platform. The notations 5.4 and 6.1 were introduced to indicate operating systems release levels V5R4 and V6R1.

The latest release of IBM i is 7.3, which delivers database, virtualization, storage management, web application serving and other enhancements that allow clients to reduce cost, improve service and manage risk of their IT infrastructure.

As a turnkey operating system, IBM i runs on a unified server platform, IBM Power Systems, which can house different logical partitions (LPAR). Each LPAR can then run different versions of IBM i. Normally, companies purchase IBM i Standard Edition with the DB2 capability built in, but IBM also offers an IBM i Express Edition option for companies deploying workloads that do not use the DB2 database.

In addition to handling data stored in the DB2 relational database, IBM i has an integrated file system that supports storage management of files in a similar way to Microsoft Windows and UNIX operating systems. The integrated file system provides a hierarchical directory structure and management interface to 11 different file systems (including UNIX, Microsoft Windows and NFS), each with its own set of logical structures and rules for interacting with information in storage. Compared to systems that focus only on their own native file system technology, the integrated file system gives companies much broader flexibility to integrate with a range of open applications from a variety of operating environments.

IBM i TERMINOLOGY

IBM i. IBM operating system. The newest version is 7.3.

Power Systems. IBM hardware supported by IBM i. This hardware is often partitioned to different logical partitions so that different operating systems can run concurrently. RAM and disk is considered part of a single storage pool in an IBM i system and is shared among different logical partitions. IBM i can share I/O feature codes such as network and Fibre Channel (FC) adapters.

Logical Partition (LPAR). Multiple workloads managed in independent operating system images. Power Systems can be divided into LPARs running different operating systems concurrently. The current operating systems supported in a Power System are IBM i, AIX and Linux.

Feature code. IBM I/O cards are sometimes known as feature codes. IBM i has FC feature codes dedicated to tape and disk. When connecting to a tape device (physical or virtual), a tape-based FC feature code is required.

Initial Program Load (IPL). The booting process for IBM i. Normal IPL default is to boot from disk, but IPL can be interrupted to boot from an alternative source including tape devices (virtual or physical). This is known as D-IPL or Alternative IPL. IBM i D-IPL functionality is supported by the DSI VTL.



Save/Restore. IBM i has a built-in backup and restore program. The native program provides the ability to write to standalone tape drives and tape libraries. GO SAV and option 21 are often used to perform a full system backup. D-IPL can be used to execute a boot from this backup. All functionality in Save/Restore is supported by the DSI VTL.

Backup Recovery and Media Services (BRMS). Licensed backup software created by IBM that incorporates media management and tape cataloging. Save/Restore does not support media management or tape cataloging. All functionality in BRMS is supported by the DSI VTL.

DSI Tracker. DSI's internally developed backup software for IBMi, Tracker provides media management, tape catalog and easy object level restore functionality. When combined with the DSI VTL, a unified backup solution is born!


DSI Conductor. DSI's VTL Conductor integrates into your existing environment and Media Management Software (MMS), delivering automated, host-based policies that leverage all of the Virtual Tape Library's (VTL) functionality without leaving the standard IBM i interface. VTL Conductor works alongside IBM's BRMS or DSI Tracker to automate the process of media duplication and/or archival activities, manage virtual resources, automatically create optimal configurations and media equilibrium. VTL Conductor can be used to manage an enterprise from single server (when BRMS is in a network), as a stand-alone product on each server, or in any combination of these modes.

IBM i ARCHITECTURE

The IBM i operating environment demonstrates a unique design that delivers innovative technology without complexity. It is a highly integrated, reliable server platform that allows businesses to run multiple operating environments. IBM i adjusts to the changing requirements of a dynamic business marketplace. IBM i offers an integrated architecture combined with high availability and security levels, simplified management, and mainframe-class technology.

One of the key contributing factors to the efficiency of IBM i is the ability to run multiple business processes and applications reliably and securely. In a study of large enterprises using multiple operating systems, IBM found utilization rates on IBM i-based servers were over 10 times higher than Intel processor-based servers and over twice as high as UNIX and other midrange-based systems (IBM i Data Sheet: IBM i = Efficient, resilient business processing, 2010). The high rate of utilization of Power Systems with IBM i is achieved through the use of a variety of proven virtualization technologies, such as sub-systems (multiple workloads managed in a single operating system image) and IBM PowerVM LPARs.

IBM i ARCHITECTURE





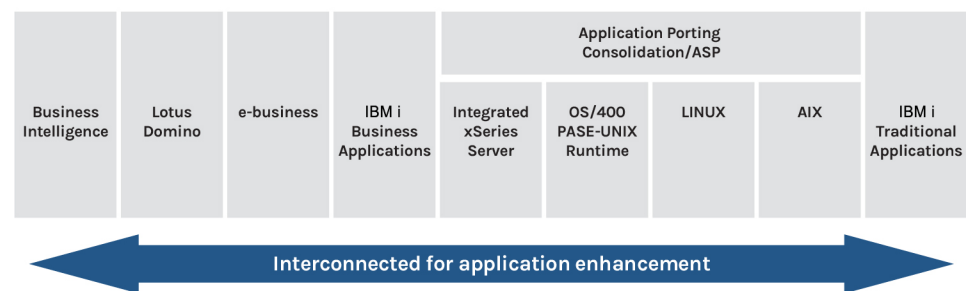
Dynamic Solutions

INTERNATIONAL

By utilizing LPARs, companies have both the power and flexibility to address multiple system requirements in a single machine to achieve server consolidation, business unit consolidation, a mix of production and test environments, and integrated clusters. LPARs are ideal for companies that want to run varied workloads in a single IBM i system. Licenses can be managed across partitions.

Each partition's system name is distinct, and the system values can be set independently. Each partition can have a different primary and secondary national language and can be operated using different time zones. This flexibility is ideal for multinational companies that wish to centralize operations in a single location while retaining the national characteristics of each system.

IBM i LPAR



IBM i BACKUP AND RESTORE

Although the IBM i is a reliable system, there are certain risks associated with backup. Data that changes often should be backed up daily. System data and data that only changes periodically should be backed up weekly. Data recovered from a site failure, disk failure, or corrupted data cannot be properly restored if it is not saved on a regular and disciplined basis.

To define a backup strategy, IBM defines the save window as the amount of time that a system is permitted to be unavailable to users while administrators perform save operations. Backup strategies are defined as simple, medium, or complex.

Simple save strategy. The simple save strategy is used by customers with a long save window such as an 8- to 12-hour block of time available daily with no system activity (including batch work). The simplest save strategy is to save everything every night or during off-shift hours. IBM i has a built-in management option 21 (entire system) that can be scheduled to automatically do these backups.

Medium save strategy. A medium save window means that there is a 4-to-6-hour block of time available daily with no system activity. Customers may use this strategy if large batch jobs are run on the system at night or large files exceed the save window.

The following techniques are used in a medium save strategy:

- Saving changed objects.
- Journaling objects and saving the journal receivers.

Save Changed Objects (SAVCHGOBJ) command is used when programs and data files are in the same library. SAVCHGOBJ saves only those files that change.

The save Document Library Object (SAVDLO) command is used to save only the documents and folders that changed. The Save (SAV) command is used to save the objects in directories that have changed since a specific point in time.



If save operations for integrated file system objects and data areas are taking too long, a journal of the objects can be used to make operations more efficient. In files with many

records, if a single record changes in a file with many records, the SAVCHGOBJ command saves the entire file.

Journaling database files and saving journal receivers regularly is a better solution, even though recovery is more complex. A similar principle applies to the integrated file system objects and data areas. Journal objects write a copy of every change in the object to a journal receiver, saving only the changed portions of the object rather than the entire object.

Complex save strategy. In instances where there is little or no time to take systems down for interactive or batch work, the complex save strategy must be employed. BRMS or DSI Tracker are typically implemented in these environments.

BRMS is IBM's backup and recovery product for the IBM i. Its policies provide the interface to manage and execute archive, backup, migration, recovery, and other removable-media-related operations. Backups using BRMS are implemented through Control Groups or native BRMS backup commands. Control Groups are analogous to Control Language Programs without the need for programming expertise. They are easier to create/customize/change without having to compile or recompile. BRMS can save libraries, objects, spool files, document library objects, and objects in directories (IFS) within the same control group. Spool files saved using BRMS can be restored with their attributes and overlays intact. Much like BRMS, DSI has created its own backup product for IBM i that has much of these capabilities in an entry-level Media Management Service called DSI Tracker.

COMMON CHALLENGES

Ensuring data is protected at remote offices is one of the many challenges IT administrators face in an edge-to-core (remote office to data center) enterprise architecture. In an environment of reduced IT budgets and resources, combined with abundant data growth, they can no longer rely on writing to tape at the remote office. In addition, organizations including those running IBM i experience the following backup and data protection challenges:

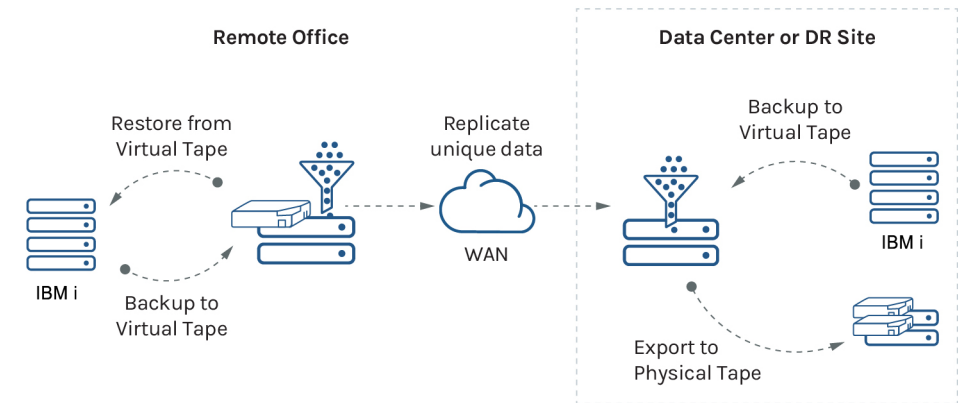
- Overwhelmed by data growth and unable to meet backup/protection windows
- Struggling to meet defined SLAs for recovery time objective (RTO) and recovery point objective (RPO) requirements
- Trying to cut down or eliminate use of physical tape and the associated storage and transport costs
- Experiencing the need to retain more backup data online without significantly expanding storage footprint
- Aging physical tape infrastructure is too slow and subject to frequent failures
- Consumption, shipment, and management of physical tape is too costly
- Desire to have multi-threaded backup targets (more LTO drives) to reduce backup queue
- Regulatory or compliance requirement to encrypt data written to tape

DSI VTL FOR IBM i

The DSI VTL is designed to increase the speed and reliability of IBM i backup tools and methods by leveraging disk to emulate all IBM i supported tape drive and tape libraries. The DSI VTL leverages existing Fibre Channel (FC) to transfer data to and restore data from a disk-based virtual tape at ultra-high speeds.



OPTIMIZED REPLICATION FOR IBM i



By using the DSI VTL with IBM i, smaller offices and remote sites can electronically vault tapes across any distance using standard IP (WAN) networks, centralizing data at the DR site for long-term data protection, retention and central tape creation. Data over the wire is encrypted (AES 256-bit) and written to the DSI VTL directly. Physical tape can also be encrypted (AES 256-bit) during export, ensuring data is safe at all times.

Because the DSI VTL emulates over 60 popular tape libraries and 50 tape-drive formats in addition to IBM i specific formats, it fits easily and transparently into the existing IBM i backup environment. Numerous tests have validated that the DSI VTL operates with all major backup software packages across multiple operating systems and environments. This allows the DSI VTL to integrate and interoperate easily with open systems backups as well as IBM i and mainframe systems, allowing consolidation of all backup operations on a single platform. With the end of ProtectTier, DSI VTL is the only qualified VTL solution for use with IBM i. Extensive testing with IBM Lab Services gives our customers the assurance they deserve.

Because data centers require both disk and tape to meet corporate or regulatory requirements for tiered backup, archive and compliance, the DSI VTL seamlessly bridges disk and tape through direct write-to-tape and best-of-breed tape management capabilities, removing additional network and operational overhead. Media management capabilities are optimized to reduce complex IT operations, reduce media consumption, and improve security through features that include:









- **Tape caching.** Transparently and directly move data from virtual to physical tape, either concurrently with backup or based on user-defined policies and schedules without the need to transfer data back across networks or through backup or media servers. Physical tape can also be migrated directly and transparently into the DSI VTL, natively and without interfering with backup operations.
- **Tape consolidation.** Write multiple virtual tapes to a single physical tape of the same or greater capacity, maximizing physical media utilization.
- **Multi-tape export.** Create multiple copies of physical tapes to meet offsite DR needs, SLAs, and regulatory requirements.

Since data security is only as good as its weakest link, securing data both on backup media and during any replication process is critical to the integrity of corporate information. The DSI VTL delivers unique features that provide comprehensive data integrity for securing data storage and migration.



- **Secure tape export via encryption.** To prevent unauthorized access to data on physical tapes, the DSI VTL solution can write to tape in an encrypted format based on the Advanced Encryption Standard (AES 256-bit). Because encryption is handled by the DSI VTL, it does not require any specific backup software or tape drive encryption licensing, nor does it require the BRMS encryption option. It works with any format or any physical tape media, maximizing media investments. The entire tape is encrypted, not just the data, allowing the physical tape to be fully protected in the event the tape is lost or stolen. For mobility, key management is included, and key packages can be exported for use with the remote VTL at the data center or DR site, allowing encrypted tapes to be imported and restored.
- **Tape shredding.** Tape shredding enables users to destroy a virtual tape image so that it cannot be accessed, even when using disk forensics. Shredding performs a three-pass wipe of the selected virtual tapes using an algorithm specified by the U.S. Department of Defense (Standard 5220.22-M).

ENCRYPTION METHODOLOGY COMPARISONS

IBM i Tape Encryption	DSI VTL Tape Encryption
 Need to deploy additional software with host-CPU impact	 No host software required, no impact on CPU
 Operational learning curve impacts staff	 Minimal learning curve, easily managed
 Limited tape drive support: TS1120 and LTO-4 only	 Encryption works with any tape drive format supported by DSI VTL
 Cannot use with existing media investment	 Utilize existing media investment: LTO-1, LTO-2, 3580, 3590, etc.

INTEGRATED DEDUPLICATION

Data deduplication dramatically reduces the amount of backup data stored by eliminating redundant data, maximizing storage utilization while allowing IT to retain more nearline backup data for longer periods and faster restores. The DSI VTL provides integrated policy-based deduplication, offering flexible post-process or inline deduplication methods, configurable for maximum flexibility and performance.

Because IBM i backups write data sequentially to tape, the DSI VTL is optimized to provide the best results based on backup and tape format awareness. A deduplication parser recognizes each IBM i tape format as well as multi-streamed backup sessions. This allows for proper alignment of the data prior to the start of the deduplication process. The metadata portion of the backup is identified and compressed. The parser is then aligned to the start of the data portion where a hash is calculated based on the best block size, and duplicate blocks are discarded. This patent-pending, tape format-aware deduplication model analyzes tape formats, guarantees that the same files are aligned the same way each time, and achieves the most efficient deduplication ratio. It allows the deduplication parser to align on different size blocks for different IBM i tape formats to ensure maximum detection of duplicate data, improving duplicate data detection by as much as 30% to 40% over generic raw fixed block deduplication analysis.



Equally important is the ability to restore data quickly from the VTL cache, physical tape, or directly from the deduplication repository. The DSI VTL is optimized for read performance, ensuring a read is nearly as fast as a write by enabling high-performance read-ahead data access, parallel read across a wide disk stripe and direct block-level access with no file system overhead. This ensures that restore speeds are fast, allowing organizations to bring critical systems back online quickly. When a file read request is initiated, the deduplication index detects the links and reads the blocks directly, in parallel from the deduplication repository, sending data blocks directly to the application for high-speed restore.

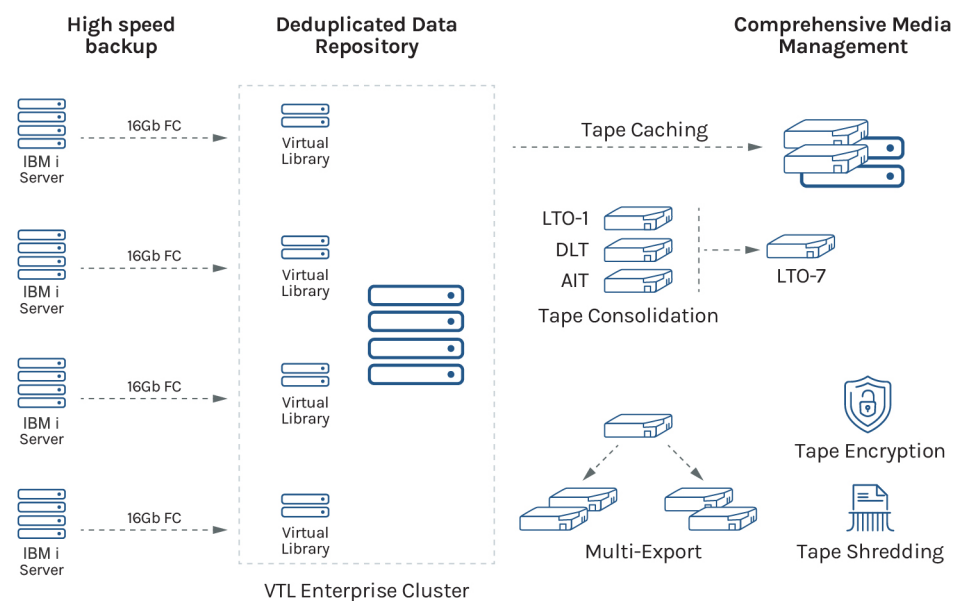
THE DSI VTL ARCHITECTURE

The DSI VTL architecture consists of three major components:

- The front-end, which is made up of virtual tape drives within virtual tape libraries using virtual tape cartridges
- A storage appliance, which is made up of an internal disk buffer, VTL and deduplication nodes
- Back-end tape drives within physical libraries using physical tape cartridges

The DSI VTL runs on Linux and emulates over 50 tape drives and 60 tape libraries provisioned via FC to the IBM i . The single-node DSI VTL server can fully saturate four 16-Gbps FC interfaces and deliver more than 6400MB/sec throughput, making it the fastest single-node performance of any VTL in the industry.

DSI VTL ARCHITECTURE




The virtual tape drives are presented to the IBM i hosts via the front-end FC interfaces and the back-end interfaces communicate with the deduplication repository to retain data on high-performance RAID 6 protected disk arrays for fast recovery and restore. The back-end interface can also attach to physical tape drives/libraries for export of VTL tapes for



long-term archival. This tape export feature supports all media management functions and has the ability to bridge normal IBM i supported physical tape by writing IBM i virtual tapes to any physical tape drives, even those that are not supported by IBM i.

DSI provides multiple deployment options for IBM i environments, including a fully integrated turnkey appliance, DSI520 Series appliances, which includes a preconfigured VTL cache and an integrated deduplication repository. For larger capacity requirements, DSI also offers a fully integrated multi-node cluster solution that scales (up to 8) for high performance and high availability combined with a shared global deduplication repository that can scale up to 4 nodes with as much as 12PB of useable RAID 6 storage deduplication repository.

DSI VTL HARDWARE CONFIGURATION OPTIONS




DSI VTL: BENEFITS TO IBM i ENVIRONMENTS

- Non-disruptive operation; IBM i does not see any difference between a physical tape and virtual tape
- Native emulation of all IBM i tape drives and libraries, including IBM 3580, 3590, 3592, TS1120 tape drives; LTO-1 to LTO-7 tape formats, and IBM 3583, 3584, and 3590 libraries
- Supports all IBM i backup tools and methods (BRMS, LXI MMS, SAVLIB, Robot Save, media policies and IPL) without any process change
- Supports mixed open system, IBM i and mainframe environments
- Ensures business continuity through IPL, bare-metal boot from virtual or physical tape
- Offers the flexibility to create as many tape targets and libraries as needed
- Higher performance (write and restore) with connectivity across 16Gb/sec FC links
- Integrated deduplication to retain backup data on capacity-optimized disk for longer periods
- Eliminates tape shipments with integrated replication and secure data encryption over the WAN
- Only unique blocks of data are transmitted, reducing WAN requirements by as much as 99%
- Direct export of virtual tape to physical tape, simultaneous with backup or based on defined policies and schedules
- The DSI VTL can cross tape-drive boundaries. Support for tape stacking allows multiple smaller virtual tapes to be consolidated onto larger physical tape media.




- Supports parallel and concurrent save/restore operations for enhanced performance:
 - **Parallel Save.** A single object is saved/restored using multiple virtual tape targets and the same IBM i job
 - **Concurrent Save.** Different objects are saved/restored using multiple virtual tape targets and different IBM i jobs

SAVE/RESTORE OPTIONS



- Business continuity; support for Initial Program Load (IPL) Bare Metal Boot from VTL or physical tape:
 - Virtual or physical tapes function as a data source for system IPL (boot)
 - TheDSI VTL provides fast recovery of systems either local or at remote site (for DR, tests, lab, etc.)

BUSINESS CONTINUITY, IPL SUPPORT



INTEGRATING THE DSI VTL WITH IBM I

The DSI VTL changes the typical approach to coordinating, managing and handling backups in an IBM i environment. Its non-disruptive approach resolves many of the issues current IBM i customers have using standard backup architectures. Implementing the DSI VTL eliminates the contention for physical tape drives, removes media and mechanical failures associated with physical tape drives and reduces the cost of physical tape transport for vaulting.

Backing up data with IBM i to a DSI VTL appliance follows the same rules and procedures as standard backup to tape. The process is completely transparent to IBM i operations, backup process, schedules, and architecture.




Using a standard FC protocol, the DSI VTL presents a virtual IBM tape library and IBM tape devices. The virtual library can accept more data streams than a physical library, and can run more jobs simultaneously. With a virtual library, the administrator can create as many tape libraries and drives as needed (up to 128 libraries and 1024 drives per node). This means that organizations can deploy multiple backup targets across IBM i systems without time constraints.

The following libraries and tape devices are certified devices to IBM i and are supported by the DSI VTL. Using these common and certified devices makes it easy for IBM i backup administrators to manage their backup infrastructure:


- IBM 3580, 3590, 3592, TS1120 tape drives
- LTO-1, LTO-2, LTO-3, LTO-4, LTO-5, LTO-6
- And LTO-7 tape formats
- IBM 3583, 3584, and 3590 libraries

ENHANCING THE BACKUP MODE

Current Model



VTL Model



One of the bottlenecks of physical tape backup is the fixed number of tape drives. Once all drives are in use, backup streams/jobs begin to queue creating a bottleneck and slowing the overall backup process. The only solution is to add more physical drives or libraries, or to resort to tape-drive sharing, which can be complex to implement, will impact performance, and could possibly require additional licensing fees.

With the DSI VTL, once the initial number of virtual tape drives is used, the administrator can simply go into the DSI VTL console and create new drives, up to the limit of the emulated library. Once a library's maximum drive count is allocated, the administrator can simply create another library on the fly.


Virtual tape replication lets the DSI VTL electronically vault tapes across any distance over standard IP networks (WAN). Because only deduplicated data is sent over the wire, bandwidth costs are reduced by as much as 90% or more.



As noted, the DSI VTL supports direct export of virtual tapes to physical tape. Because data is streamed directly from the DSI VTL to tape, tape drives can stream at their maximum performance levels, speeding the tape creation process and reducing the number of tape drives needed to produce media.

Normally, IBM i recognizes only a few physical tape formats. By leveraging the DSI VTL you can bridge these IBM i virtual tapes to any physical tape, including those not supported by IBM i systems. For example, an LTO-2 virtual tape can be copied to an STK 10000 physical tape. Tape stacking consolidates multiple smaller virtual tapes into a larger physical tape format, maximizing media utilization while reducing tape consumption and cost.

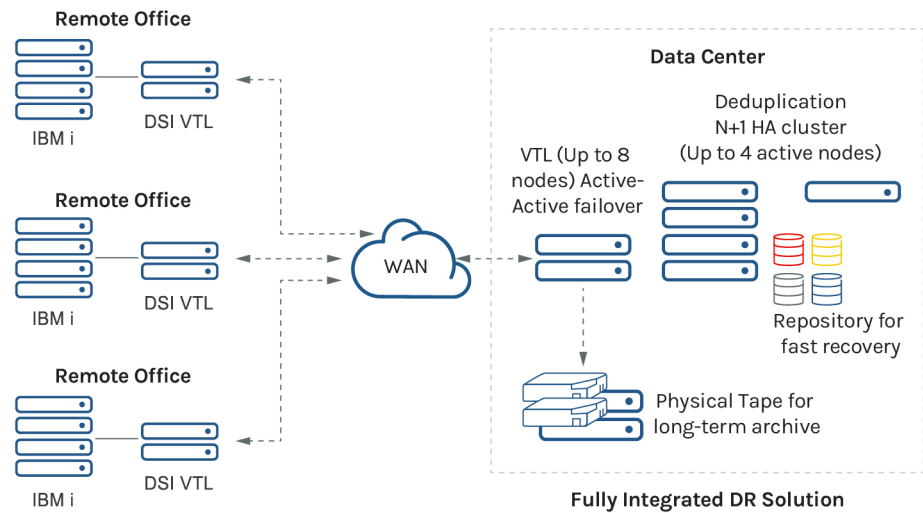
ELECTRONIC VAULTING



Built in replication is used for tape vaulting and DR, eliminating tape transport and providing the opportunity to create secure physical tape at the DR site. Running wizard-based policies, each individual virtual tape or job can be deduplicated, compressed, replicated, and retained at the DR site in the online disk repository for a defined period of time, and then vaulted to physical tape with a unique policy and a defined schedule. Facilitating replication offers additional cost savings by reducing WAN bandwidth requirements, eliminating the need to create tape at the remote or branch office, and the associated tape transport cost while ensuring security. Added benefits include support for IPL (boot); administrators can D-IPL from the virtual tape or from physical tape, locally or from the DR site, ensuring fast recovery in the event of a catastrophic failure. Additionally, the DSI VTL encryption configuration simplifies the process of selecting tapes to encrypt, and determines where to store the tapes and associated secure encryption keys. The process is transparently managed and controlled by the DSI VTL, offloading the computing resource required by IBM i software encryption as well as the need to procure hardware encryption devices to meet security requirements.



FULLY INTEGRATED DR SOLUTION



CONFIGURATION GUIDELINES

The DSI VTL is available in several form factors, including an integrated turnkey appliance, a VTL node using SAN storage, or as a high availability (HA) cluster, which provides the added benefit of higher performance, failover and higher scalability for longer term data retention.

Storage Appliance. The DSI VTL storage appliance is an all-in-one appliance that includes all licenses for virtual tape management, deduplication, replication, physical tape management, and encryption features. The DSI VTL console provides two methods for creating VTLs: a configuration wizard, or the ability to select the Virtual Tape Libraries object and create a new library from a predefined list. Once the administrator creates the emulated tape library and drives, he or she will be prompted to create virtual tapes. At this point the new VTL is assigned to the IBM i SAN client, and all existing backup tools and processes will operate without change. For more details on the easy-to-use process, please reference the latest DSI VTL User Guide.

VTL Nodes. The VTL Node operates and is licensed in a similar manner to the DSI VTL storage appliance. The primary difference is that the DSI VTL Node does not include physical disk. Instead, it attaches to an existing SAN via FC and zones storage resources from that existing SAN.

HA Cluster. The HA cluster configuration is a highly scalable HA configuration that includes independent VTL and deduplication nodes, providing flexibility within each individual



environment. Customers can add more VTL nodes for faster backup performance or add more deduplication nodes for higher deduplication performance or longer data retention. This cluster configuration can scale to 8 VTL nodes, 4+1 deduplication nodes, and up to 12PB of physical RAID 6 usable storage (deduplication repository).

Configuring and administering backup policies to the DSI VTL is the same as with a physical tape library. Once a virtual IBM tape library and drives are assigned to the IBM i, operations continue as normal.

IBM i backup administrators can start with the IBM i native backup and restore tools to back up data to the DSI VTL. As backup needs grow, the administrator can move to more sophisticated IBM backup software, such as BRMS or DSI Tracker and continue to use the DSI VTL without change. A single VTL can be partitioned to multiple IBM libraries with multiple drives. Each partition can be dedicated to different LPARs providing, parallel backup streams and reducing backup queues and wait time.

Data recovery is also significantly faster since restores are directly from the DSI VTL. Faster restore performance is the result of advanced read-ahead technology, parallel reads across a large array of disk, direct block-level access, and no file system overhead. When a read operation is performed, the VTL reads ahead sequentially on the virtual tape and adds the read ahead buffer into a queue. The implementation of this feature dramatically improves read performance as well as deduplication performance. Using this advanced technology; DSI has demonstrated that the restore speed from deduplicated data and non-deduplicated data is almost exactly the same across all systems, including single-node VTL appliances and cluster solutions.

Meanwhile, a copy of the data can be electronically vaulted to the DR site, retained on physical disk in a deduplicated format, or securely exported to a physical tape using the integrated encryption configuration tools. Administrators can restore any required data directly from the DR site on the remote DSI VTL or from physical tape.

Encrypt data on virtual and physical tapes

To ensure that the data exported to tape is confidential and secure, the DSI VTL enables encryption keys to be created that use the Advanced Encryption Standard (AES) 256-bit key algorithm (Secure Tape) published by the National Institute of Standards and Technology. When data is exported to physical tape, an encryption key is chosen to encrypt the data. That same key must be used to import the same data back to virtual tapes or to read the data.

Each key consists of a secret phrase. For additional security, each key is password protected. A single key can be applied to all virtual tapes when exported to physical tape, or a unique key can be created for each tape. Once one or more keys have been created, they can be exported to a separate file called a key package. If encrypted tapes are sent to other locations that run the DSI VTL, the key package can be sent so tapes can be decrypted remotely.

The DSI VTL includes a key management system, which provides a simple interface for creating, managing and maintaining keys used to encrypt the data. Once encrypted, the data is not accessible without access to or knowledge of the encryption key. The key management system helps users create high-quality keys that are impossible to crack, but easy to manage and maintain.

The entity that encrypts the data is the owner of the data. It is also the same entity that decrypts the data when needed. Even in the case where another entity needs to decrypt the



data, that entity must be authorized with the same level of trust. There is no unknown entity in the trust relationship. Benefits of the DSI VTL encryption include:

- Key packages can be exported for use on remote VTLs (to allow tapes to be imported/restored).
- The encryption is handled by the VTL and does not require an encryption license on IBM i or backup software.
- Host processing power is not affected by the encryption process.
- The encryption option works with deduplication and replication. This includes the replicated data (unique deduplicated data and metadata) that can also be encrypted.
- The encryption option can work in conjunction with tape stacking (tape consolidation) for long-term tape archival.
- Data is encrypted on both virtual tape and physical tape.

Sizing

Much has been written about deduplication ratios, and many vendors have a tendency to indiscriminately advertise large deduplication ratios. However, no vendor can guarantee a specific deduplication ratio without first analyzing a customer's specific data and backup methodology. The ratio of data stored (after deduplication and compression) wholly depends on a number of different factors, including:

- Backup methods and policies (full, incremental, daily, weekly, etc.)
- Disk retention policies
- Rate of data change (for example, database data and email data that changes frequently has lower deduplication ratios)
- Type of data (structured, unstructured, images, etc.)

Sizing a DSI VTL solution must take these backup methods and policies into account. Since the IBM i is not aware of deduplication, a few general rules need to be followed:

1. Perform an assessment of the backup solution so DSI VTL sizing can be done in collaboration with the backup team. A discussion of methodologies, long-term intentions, and aforementioned factors should be discussed in detail to understand the customer's environment, retention policies, SLA and DR environment. If a customer has BRMS in use today, reports can be created for use with DSI's sizing tool.
2. Plan for the initial disk design. The size of the DSI VTL physical resources are calculated by taking into consideration the size of the total backup source (all data combined) and anticipating a deduplication ratio based on data type, change rate, retention, etc. This will provide a strong phase 1 estimate of the required size.
3. It is critical to understand that the benefits of data deduplication are realized over time. Initial deduplication ratios may be relatively low during the first few weeks of backup. This is normal since larger deduplication ratios are a factor of retaining data for longer time periods, combined with lower rates of changed data. These ratios will improve over time. DSI recommends recalculating the DSI VTL sizing once a significant subset of production backup data is received from an IBM i. This will provide a more effective view of long-term deduplication ratios and a better plan for future growth.

By adopting an incremental approach, a backup administrator will be better able to forecast the space needed to sustain data growth. This method requires ongoing measurement that will help the administrator grow the system over time.



SUMMARY

Data is driving our world; information is becoming the new currency. While technology changes, data remains irreplaceable and a corporation's most valuable asset. But protecting that asset presents many challenges. IT budgets are flat or down, and data continues to grow at a rate of 50% to 60% annually.

Most IBM i installations continue to use tape, stockpiling and trucking data to tape silo centers. Furthermore, corporations are saving information for longer periods of time to meet regulatory and compliance requirements, and data is becoming more difficult to find and restore in a timely fashion. IBM i backup administrators face these challenges daily.

Adapting the DSI VTL into the IBM i infrastructure allows customers to take full advantage of the product's deduplication, replication and tape media management with encryption, with little or no disruption to backup operations. The DSI VTL solution is an ideal backup-to-disk target for IBM i backup environments, allowing organizations and their administrators to:

- Reduce backup storage requirements by as much as 95% through deduplication
- Meet defined SLAs for RTO and RPO requirements by retaining data on disk for longer time periods
- Reduce/eliminate tape expenses by extending disk retention policies
- Eliminate offsite tape management and transportation costs using secure electronic vaulting
- Meet regulatory and compliance requirements through encryption
- Improve resource management through a single platform for open systems, IBM i, and mainframe

For more information, visit www.dynamicsolutions.com or contact a DSI team member.

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 Software. All rights reserved.