



# QuickRecover™

---

Cost-effective Data Protection for your Mainframe

*The Shoden QuickRecover™ solution is designed to meet the needs of all mainframe users wanting to improve resilience and data protection while cutting the costs of using physical tape for backup. Shoden Data Systems is committed to delivering high-end architectures and solutions that ensure unsurpassed data availability, scalability, and industry-leading performance to leading corporate clients. These solutions are based on core technologies from organisations such as: Hitachi, Data Domain and Luminex.*

## **QuickRecover™ – the new way to reduce costs, while improving resilience and protecting data**

### **INTRODUCTION**

Many mainframe sites are now looking to retire their current tape infrastructures and to eliminate their dependence on physical tape and the associated movement of physical tape for off-site vaulting and recovery. Virtual Tape Libraries have become common in many enterprise sites due to the problematic nature of physical tape media and the integrity of the data placed on them. Data growth has also resulted in longer backup times and shortening backup windows. Governance requirements have also taken their toll on backup environments by enforcing strict adherence to data management and off-site storage regulations.

### **TRADITIONAL BACKUP ENVIRONMENTS**

Until recently, disks have been economically unsuited for backup and recovery storage. Disk-based data protection methods such as replication were once reserved exclusively for critical data with the highest availability and recovery requirements. Even that data has historically been backed up to tape. Moreover, to provide a historical repository and to recover from user or operator errors, such as accidental deletion of files, multiple versions of data must be backed up over time, magnifying the cost problem. With rapidly falling disk prices and cost-effective Fiber Channel based disk arrays, disk-based storage and network replication now offer a new tool for protecting data at an affordable price point.

However, simply inserting cheap disk into the process does not provide better data protection or optimise the total cost of backup.

#### *Tape: Advantages and Challenges*

For several decades, tape has been the medium for storing backup data and for transporting replicated data offsite for disaster protection. With increasing capacity and performance, tape allows storing multiple copies or versions at a low cost.

Unfortunately, tape also has some drawbacks – it is optimized for backup, but not for restore. With the increasing transfer rates of tape drives, backups need to be carefully staged to stream the tape drive to avoid the “shoe shine” effect of starting, stopping, and repositioning the tape. Incremental backups only worsen the problem since incremental backups do not generate high enough data rates. Thus, streaming a tape drive typically requires “multiplexing,” the blending of concurrent backup streams from multiple clients. By maximizing tape drive utilization, multiplexing helps backup performance, but slows down restore performance because of the need to skip data belonging to other backup clients.

However, the biggest challenge for tape is uncertain data integrity. A backup process may have completed successfully, but verifying the data on all of the tapes is effectively impossible without doing an actual restore. Surveys have shown that anything up to a third of all backups fail to restore successfully. One bad tape can cause a restore operation to fail and render the entire series of tape media useless, and often is not discovered until an actual restore operation is under way.

#### *Disk: Advantages and Challenges*

Disk storage offers several advantages over tape. Unlike tape drives, disk arrays do not need a steady stream of data. Even incremental backups that generate small amounts of data do not create a “shoe shine” effect.

---

Secondly, disk arrays can simplify and speed up the overall backup process by allowing the administrator to perform fewer full backups without suffering a performance penalty, or increasing the risk of failing restores. Despite shrinking backup windows when using tape, frequent full backups are performed to minimize the number of tapes required for restores, simplifying the recovery process. The number of tapes required for a restore increases with incremental backups, which increases the time required for restore and the risk that one of the tapes is unrecoverable.

Disk allows the administrator to shorten backup windows. Disk also makes off-site recovery copies easier and more efficient. Most importantly, disk is superior for recovery – in reliability and performance. Disk-specific technologies like RAID make disk a more reliable medium than tape. As mentioned, one bad tape in a sequence can cause an entire restore operation to fail. With RAID protection, a restore can continue and complete successfully even with a failed disk. (Whilst vastly better than tape, RAID is not by itself sufficient to guard against all issues, as explained below.)

Moreover, according to Strategic Research, 87% of all restores are single file recoveries, not full system recoveries. As a random access device, disk enables much faster single file recoveries. Average access times for disk are measured in milliseconds. In contrast, average file access times for tape, a serial access device, range from 27 to 73 seconds. If a restore requires tape location, loading, and unloading operations, the overall “time-to-data” is even worse.

Despite the advantages, the biggest limitation for standard disk storage for backup retention has been cost. Even with dramatically falling disk prices, tape has been the most economic choice. That is why it is still around. The cost equation is simple. Imagine storing four weeks of weekly full and daily incremental backups. Assuming incremental backup sizes are approximately 5% of the original data, the storage requirement is five times the original data size (4 weekly + 0.05x20 incrementals = 5).


### THE SHODEN QUICKRECOVER™ SOLUTION FOR MAINFRAME BACKUPS

The challenge is clear. A backup storage solution should combine the economy of tape with the usability and speed of disk. At the same time, it must overcome the disadvantages of both tape and conventional disk storage arrays. Special requirements for backup storage are:

- ◆ **Economy**
- ◆ **High Performance**
- ◆ **Data Integrity**
- ◆ **Minimal Disruption**

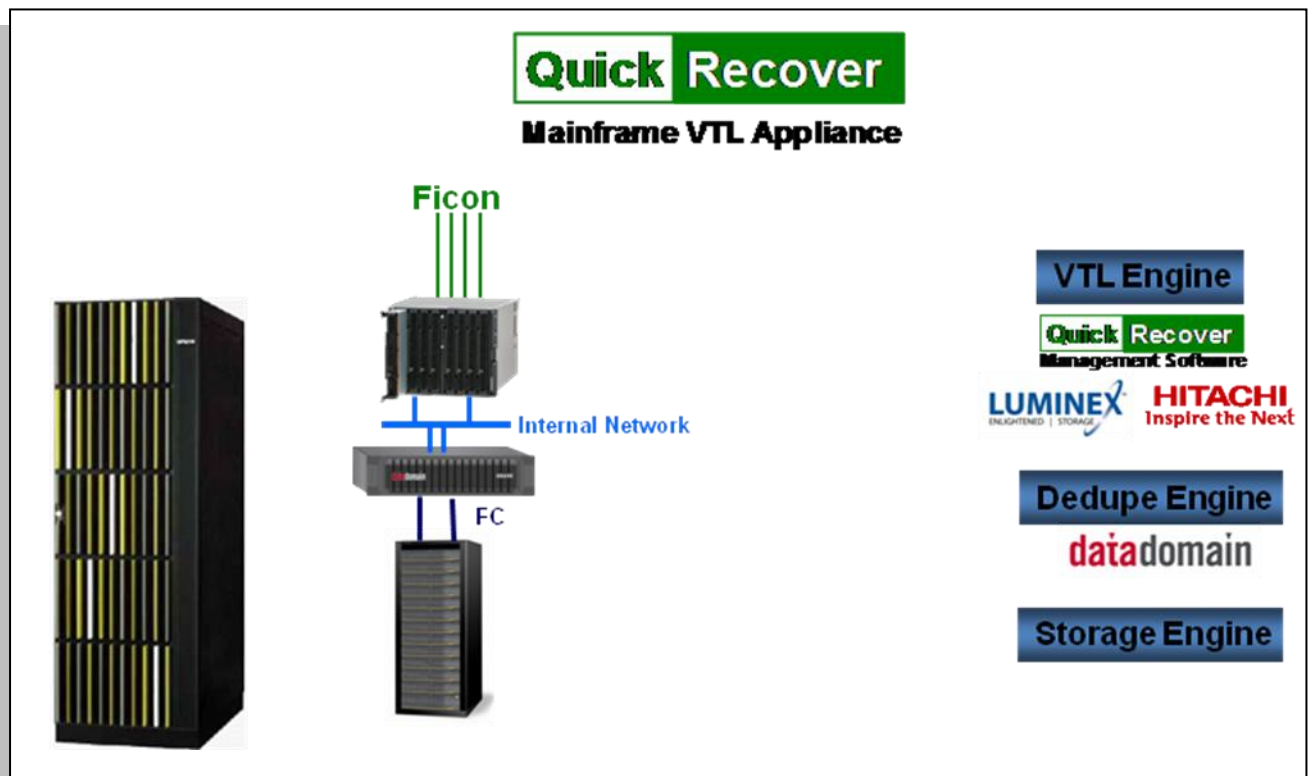
To overcome the issues related to backup operations, governance requirements and recoverability of tape media, Shoden Data Systems has developed the “**QuickRecover™**” **Mainframe VTL Appliance**. Using an integrated system of components that are internationally recognized, the unique disk-based offering brings de-duplication into the Mainframe environment and allows the reduction or removal of tape dependant backup, recovery and vaulting.

The "QuickRecover™" Mainframe VTL Appliance is built on proven technologies that integrate to form a complete solution.



The QuickRecover™ appliance contains four key elements, all contained within one frame.

- The VTL engine, Hitachi Blade Servers, together with Luminex's Channel Gateway.
- The QuickRecover™ management software, running on one of the blade servers, controlling, monitoring and reporting on all aspects of the appliance
- The Dedupe engine, powered by Data Domain to compress and de-duplicate the backup data, and enable it to be stored on disk, and optionally to provide offsite replication
- The Storage Engine, to store the de-duplicated backup data.



### *The Luminex Channel Gateway*

Luminex's Channel Gateway enables open systems disk storage solutions to emulate a mainframe tape library and enables the mainframe enterprise to exploit the latest features being delivered by the disk storage industry. Simply put – the Channel Gateway eliminates the proprietary boundaries and vendor lock-ins that separate the mainframe from the open systems storage world. The Channel Gateway has been tested in the SNIA/SNW laboratories and simulates both 3480 and 3490 mainframe tape Drives.

Keeping all production data on-line is now a realistic possibility. With the Channel Gateway, organisations can now use high capacity, cost effective, open systems storage subsystems as replacements for their mainframe tape libraries. Rather than being forced to consider mainframe tape and library solutions from a few proprietary vendors, enterprise managers can now consider solutions from the entire open systems storage industry. The capacity, functionality and cost advantages of the Channel Gateway open up a new world of options for the mainframe community. Now one can realistically and cost-effectively hold many terabytes of mainframe backup data online.

### *The Data Domain Enterprise Series*

The Data Domain Enterprise Series was designed from the start to store backup data and to enable automated offsite replication, all driven by standard enterprise backup software. As high performance, online backup appliances, Data Domain Appliances enable faster backups, restores, and replication to meet shrinking backup windows and rapid recovery requirements. Restorers integrate seamlessly with your existing backup/recovery processes, architectures, and replication to offsite storage, fully leveraging investments in backup software and training.

What makes Data Domain unique is its operating systems software. Data Domain OS offers unprecedented levels of protection, verifiability and self-healing capabilities, through its adoption of Data Invulnerability Architecture, unavailable in conventional disk or tape systems. And with its unique de-duplication technology, the Data Domain OS can store several months of recovery copies in an extremely small number of disk drives, lowering the price/GB and greatly simplifying administration.

Data Domain solutions are designed to meet the unique demands of backup and recovery storage.

- ◆ **High performance for both backup and restore:**

Unlike a tape drive, a Data Domain does not require constant streams of data for the best backup performance. Multiple backup streams at varying speeds can be sent to one Data Domain. A Data Domain appliance also enables fast single file restores by taking advantage of the random access nature of disk.

- ◆ **Economy:**

The Data Domain OS dramatically reduces the storage required for backup data by pooling redundancies within backup images and storing only unique data patterns. This allows the Data Domain to not only detect and eliminate storage of duplicate files but also to detect and eliminate repeated patterns within and across files. With its unique Global Compression technology, the Data Domain OS delivers an effective compression ratio of up to 20:1 over time. As a result, a Data Domain Restorer is an order of magnitude smaller, simpler and easier to administer than cheap disk arrays used for backup storage.

◆ **Data Invulnerability:**

A Data Domain appliance is designed to prevent, detect, and heal from hardware or software failures to ensure data integrity and restorability. The Data Domain OS file system and RAID design were built from the ground up to offer fault protection, detection and correction from software flaws and disk errors in a much more rigorous way than general purpose disk storage or file systems.

◆ **Easy to use and integrate into an existing backup software environment:**

To a storage administrator, a Data Domain Appliance has a familiar “feel” with its industry standard NFS or CIFS interface. A Data Domain Appliance fits easily into an existing backup environment.

Data Domain allows large volumes of data to be virtually “compressed” onto less physical storage. This is done through two mechanisms which provide virtual capacity in excess of up to 10x that of the physical capacity on which the data is stored. (In open systems environments, larger deduplication ratios are commonly observed).


The first mechanism that Data Domain uses to achieve such high levels is compression. The first time data is sent to the Data Domain appliance, it is compressed and stored on the external storage. Compression ratios of between 2:1 and 3:1 are possible, depending on the type of data.

The second mechanism used by Data Domain is the process of de-duplicating previously “seen” data. When a backup is run for the second or subsequent time through the Data Domain appliance, the data is inspected for common data elements already stored in the Data Domain. In most environments the daily data changes make up a small percentage of the entire storage used and will remain constant over time. Data Domain understands this phenomenon and intercepts this data and references it to the data elements already stored. Any changed data will be classed as “new” data when passed through the Data Domain, and the compression algorithm is used to compress and store the data blocks.

By using these two methods the Data Domain can store far more data than the actual physical capacity, allowing for disk-based backups to be stored on the Data Domain for several years. Retention periods of the data are managed by the chosen backup product and passed through to the Data Domain, which does regular cleanups of the storage, as in the case of normal tape usage.

*Data Domain-to-Offsite Replication for Disaster Recovery*

The Data Domain approach makes off-site replication or network vaulting much easier, faster and affordable. Optional Data Domain Replicator software takes advantage of the unique Data Domain OS features to create a simple and efficient mechanism that replicates backup data asynchronously across a network between two Data Domain appliances. With Replicator, data is backed up to the local Data Domain using standard backup software. The local Data Domain then acts as an “originator” and replicates the data over a TCP/IP LAN or WAN to a remote Data Domain replica.



Thanks to the de-duplication, there is a massive reduction in the amount of data that needs to be sent over a WAN. Only changed and de-duplicated blocks are transmitted across the WAN, so with a 10:1 deduplication ratio, there is a 90% reduction in WAN traffic, when compared to replication from any other disk-based backup system. If necessary, data can be restored across the network from the remote site with the same efficiency or the data can be restored at the remote site.

The disk storage options for the Data Domain can be either Fibre Channel (FC) Disk Drives or SATA Disk Drives. Due to the nature of the data, FC Drives are generally used within QuickRecover for Disk based backup, ideally protected by RAID-5 or RAID-6 configurations.

### *Hitachi Blade Servers*

Hitachi has gained an enviable reputation throughout the I.T. industry for the robustness and resilience of its products. The Shoden QuickRecover™ solution makes best use of these characteristics with Hitachi's Blade Servers.

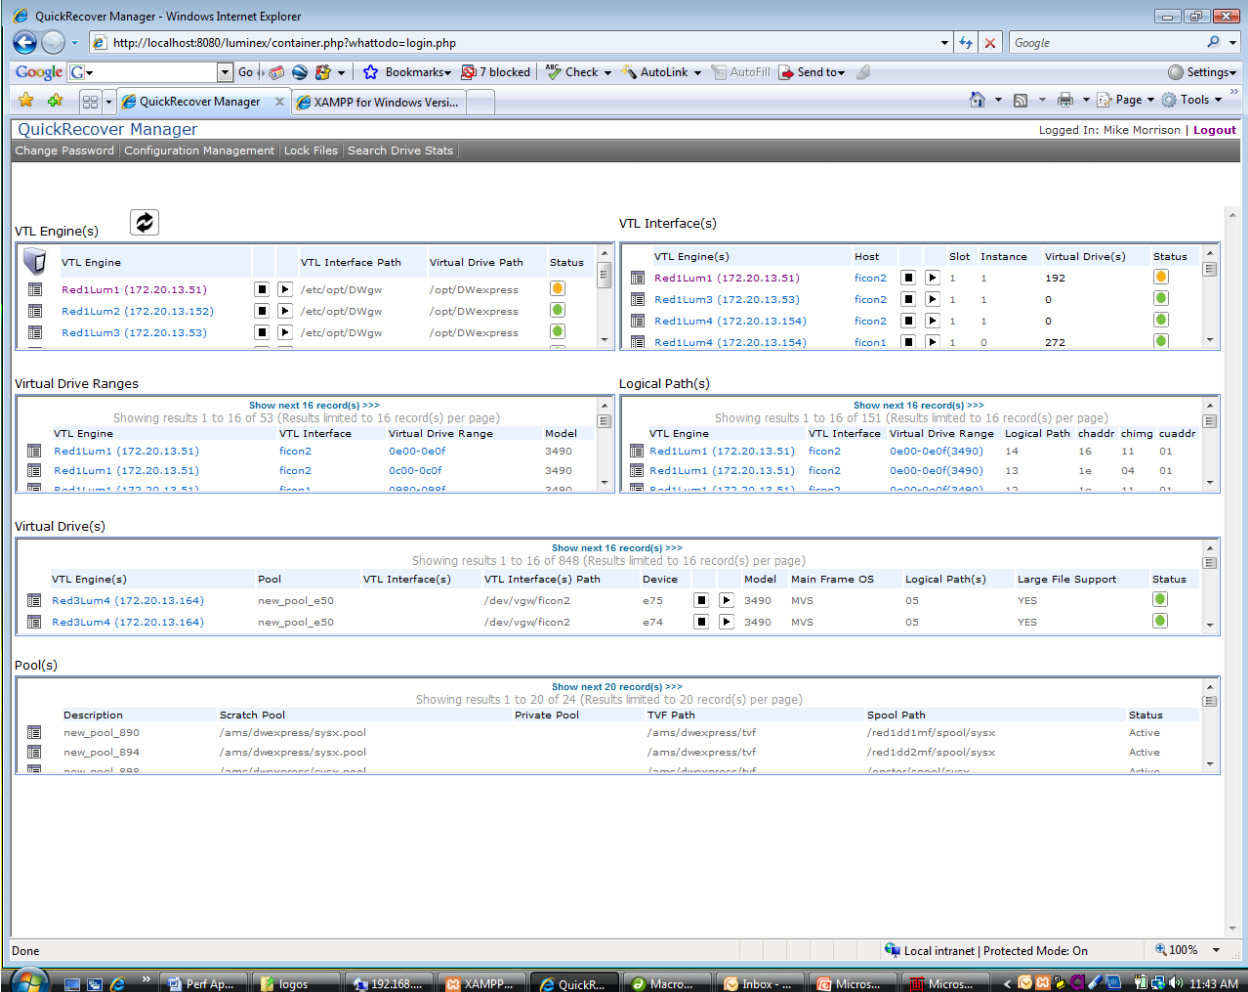
## **HITACHI** Inspire the Next



Hitachi provides the industry's first enterprise level blade server, with the capability to have both Itanium and Xeon in the same chassis, the ability to scale up to 2x16 or out to 8x8, with N+M automatic failover, and with its own embedded virtualisation.

## Shoden's QuickRecover™ Management Software

Shoden Data Systems has developed unique management software to provide a single interface to the QuickRecover™ Solution and provides the integrated configuration management and performance monitoring necessary to meet the operational requirements of the most demanding environments..



QuickRecover Manager - Windows Internet Explorer  
 http://localhost:8080/luminex/container.php?whattodo=login.php  
 Logged In: Mike Morrison | Logout

**VTL Engine(s)**

VTL Engine	VTL Interface Path	Virtual Drive Path	Status
Red1Lum1 (172.20.13.51)	/etc/opt/DWgw	/opt/DWexpress	On
Red1Lum2 (172.20.13.152)	/etc/opt/DWgw	/opt/DWexpress	On
Red1Lum3 (172.20.13.53)	/etc/opt/DWgw	/opt/DWexpress	On

**VTL Interface(s)**

VTL Engine(s)	Host	Slot	Instance	Virtual Drive(s)	Status
Red1Lum1 (172.20.13.51)	ficon2	1	1	192	On
Red1Lum3 (172.20.13.53)	ficon2	1	1	0	On
Red1Lum4 (172.20.13.154)	ficon2	1	1	0	On
Red1Lum4 (172.20.13.154)	ficon1	1	0	272	On

**Virtual Drive Ranges**

VTL Engine	VTL Interface	Virtual Drive Range	Model
Red1Lum1 (172.20.13.51)	ficon2	0e00-0e0f	3490
Red1Lum1 (172.20.13.51)	ficon2	0c00-0c0f	3490
Red1Lum1 (172.20.13.51)	ficon1	0800-080f	3490

**Logical Path(s)**

VTL Engine	VTL Interface	Virtual Drive Range	Logical Path	chaddr	chimg	cuaddr
Red1Lum1 (172.20.13.51)	ficon2	0e00-0e0f(3490)	14	16	11	01
Red1Lum1 (172.20.13.51)	ficon2	0e00-0e0f(3490)	13	1e	04	01
Red1Lum1 (172.20.13.51)	ficon2	0e00-0e0f(3490)	12	1e	11	01

**Virtual Drive(s)**

VTL Engine(s)	Pool	VTL Interface(s)	VTL Interface(s) Path	Device	Model	Main Frame OS	Logical Path(s)	Large File Support	Status
Red3Lum4 (172.20.13.164)	new_pool_e50		/dev/vgw/ficon2	e75	3490	MVS	05	YES	On
Red3Lum4 (172.20.13.164)	new_pool_e50		/dev/vgw/ficon2	e74	3490	MVS	05	YES	On

**Pool(s)**

Description	Scratch Pool	Private Pool	TVF Path	Spool Path	Status
new_pool_890	/ams/dwexpress/sysx.pool		/ams/dwexpress/tvf	/red1d1mf/spool/sysx	Active
new_pool_894	/ams/dwexpress/sysx.pool		/ams/dwexpress/tvf	/red1d2mf/spool/sysx	Active
new_pool_898	/ams/dwexpress/sysx.pool		/ams/dwexpress/tvf	/red1d3mf/spool/sysx	Active

Key elements of the QuickRecover™ management software are:

- User authentication with role-based access
- Configuration management
- Auto-Discovery of server configuration when another engine is added to the QuickRecover™ environment
- Administration of custom drive configuration by grouping ranges or individual virtual drives across the environment into separate pools or “Esoterics”
- Channel, Path and drive assignment
- Searching log statistics and messages throughout the QuickRecover™ system based on date, status, individual drives or volumes
- Facility to stop or start individual VTL Engines / Ficon Interfaces / Drives

## USER SUCCESS STORY

Shoden's QuickRecover™ solution is being successfully used by a variety of users.



Edgars Consolidated Stores, South Africa's largest retailer, employs over 20,000 people, and has a turnover of approximately \$4B. Edcon is the leading clothing, footwear and textiles (CFT) retailing group in southern Africa trading through a range of retail formats. The Company has grown from opening its first store in 1929, to ten retail brands trading in over 1000 stores in South Africa, Botswana, Namibia, Swaziland and Lesotho. In addition, Edcon operates world-class credit, information technology and distribution systems, being a major credit card provider.

Edcon's IT organisation faced the following problems with their mainframe backups:

- Floor space: the physical size of their existing tape libraries
- Growing media and drive errors
- High maintenance costs of their tape infrastructure
- No HSM ML2 data offsite
- Complexity of remote tape restores

By adopting Shoden's QuickRecover™ solution, Edcon saw the following benefits:

- Tape and media failures eliminated
- Full backups replicated to DR site
- 11:1 backup data deduplication achieved
- All former tape data now included in DR plan
- Much lower maintenance costs
- Redeployed tape staff
- Simplified and faster restores

## BENEFITS OF THE SHODEN QUICKRECOVER™ APPROACH

The QuickRecover™ Mainframe VTL Appliance will provide the following benefits:

- ✓ Disk-based storage for backup data, reducing the probability of tape media failure
- ✓ Much higher average backup success rate due to lower mechanical failure.
- ✓ Lower human intervention than with tape media
- ✓ Automated Off-Site backup data replication
- ✓ Improved restore times as the requirement for tape search is eliminated
- ✓ Highly scalable solution for unforeseen volume growth
- ✓ Redundant QuickRecover Appliance (Primary and DR Sites) in case of failure
- ✓ Data Domain compression and hash algorithms maintain performance over time, unlike other VTL products which suffer from storage fragmentation.
- ✓ Data Domain has the most experience in de-duplication technology in the industry

- ✓ The technology is proven and tested.
- ✓ Simplified upgrade path for all components of the solution

*Shoden Data Systems delivers high-end architectures and solutions, ensuring unsurpassed data availability, scalability, and industry-leading performance to leading corporate clients. These solutions are based on proven technologies. The Shoden QuickRecover™ solution meets the needs of all mainframe users wanting to cut costs while improving resilience and data protection.*

Philip Jones  
Chief Technology Officer  
Shoden Data Systems UK  
18/12/2008