



QuickRecover™ Mainframe VTL Appliance

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 Data Domain and Luminex.

QuickRecover™ – the new way to reduce backup costs in the mainframe environment while improving resilience

INTRODUCTION

Many mainframe tape infrastructures are now reaching their end of life (EOL). Organisations are looking to eliminate their dependence on physical tape and off-site vaulting of tape. The time taken to recover after a failure or disaster, is critical to some organizations' chances of long term survival; unless key applications are up and running within 1-2 hours, business losses are likely.. Governance requirements have enforced stricter adherence to data management and off-site storage regulations. Users are demanding faster recovery of their data; growth has lead to longer backup times and shortening backup windows.

Today, installations are looking to achieve the following:

- ✓ Eliminate Tape Silo maintenance and running costs
- ✓ Reduce power, cooling, floor space costs
- ✓ Automate offsite backup replication
- ✓ Unify backup across multiple platforms (z/OS, Wintel, UNIX)
- ✓ Remove offsite transport costs
- ✓ Improve restore times to meet RTOs (recovery time objectives)

TRADITIONAL BACKUP ENVIRONMENTS

Until recently, disks have been economically unsuited for backup and recovery storage. The continuous decline in the unit cost of disk storage over the last twenty years has also altered the dynamics of media selection for backup. Historically, the unit cost of tape storage has been about one tenth of that of disk. The rapid growths in disks' areal density and the adoption of deduplication techniques have changed the cost equation in favour of disk for backup. 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 Fibre 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 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, because they do not generate high enough data rates. Thus, streaming a tape drive typically requires “multiplexing,” the blending of concurrent backup streams from multiple sources. 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 the problem is not discovered until an actual restore operation is under way.

Disk: Advantages and Challenges

Using disks for backup 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 (Redundant Array of Independent Disks) 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

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, illustrated in Figure 1.



Figure 1- QuickRecover Mainframe VTL Appliance Components

The QuickRecover™ appliance contains three key elements, all contained within one frame (illustrated in Figure 2).

- The VTL engine, 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 Deduplication 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

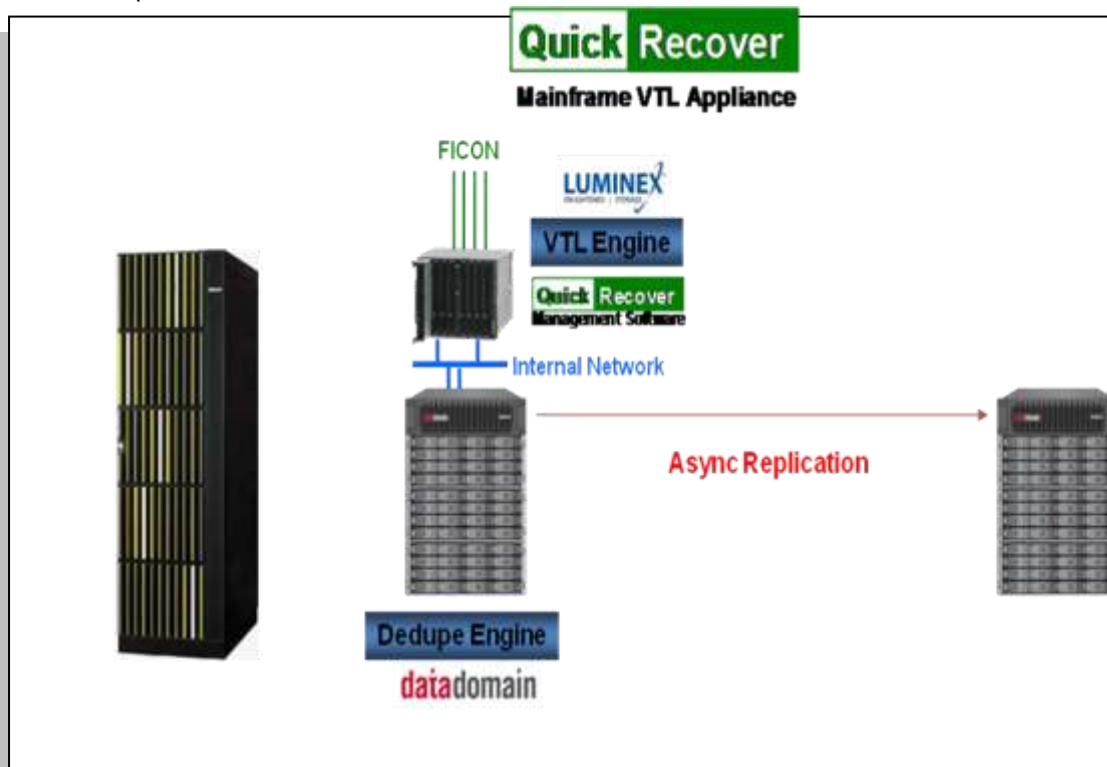


Figure 2- QuickRecover VTL Appliance Key Elements



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 (Storage Networking Industry Association) laboratories and emulates 3480, 3490 and 3590 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 Luminex Channel Gateway open up a new world of options for the mainframe community. Now, one can realistically and cost-effectively hold petabytes of mainframe backup data online.

The Luminex Channel Gateway is a virtual tape solution emulating a mainframe tape control unit that enables mainframe tape to be stored and managed, using modular, open systems disk solutions. The Luminex Channel Gateway allows you to deploy proven Fibre Channel, iSCSI and NAS storage devices from industry leading vendors. With the Channel Gateway, your existing tape backup and recovery applications continue to be used, without change. FICON and ESCON connectivity are both supported, so forced upgrades to FICON are not required. The Luminex solution simply deploys Channel Gateways, with proven Fibre Channel, iSCSI or NAS storage devices. The Channel Gateway works with your existing mainframe storage management applications (DFDSS, CA1, DFHSM, RMM, FDR etc). Introduction of the Channel Gateway into your data centre is fast and straight-forward. With the Channel Gateway, no changes to your existing backup software or processes are necessary, so you can quickly and seamlessly implement the solution for production operations.

The Channel Gateway enables you to get away from high cost proprietary and sole source solutions. It represents the greatest degree of options and flexibility. In short, the Channel Gateway is the most cost-effective mainframe storage solution both in terms of current acquisition and future recurring costs. Channel Gateway products are based on state-of-the-art, highly reliable servers. Luminex has not only tested and certified in its mainframe labs with partners and an extensive storage infrastructure, but they have also certified with the world-class testing suites from a major mainframe vendor. The products are customer proven at the largest customers worldwide. Luminex offers five-levels of customer support including 24x7 on-site, and is proud of its customer-first approach.





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 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. (There are several broadly similar compression

algorithms across the IT industry that are used to remove repeated characters and blank spaces). 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 disk. 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, illustrated in Figure 3, 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.




Figure 3- Replication for Disaster Recovery

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 (typical in mainframe environments), 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. For other environments, with higher deduplication ratios, there is an

even greater reduction in WAN traffic. Data Domain's internal storage architecture allows non-disruptive, rapid and simple capacity upgrades in a RAID-6 (double parity) protected environment.

Shoden's QuickRecover™ Management Software

Shoden Data Systems has developed unique management software to provide a single interface to the QuickRecover™ Solution, illustrated in Figure 4, providing the integrated configuration management and performance monitoring necessary to meet the operational requirements of the most demanding environments. Key elements of the QuickRecover™ management software are:

- User authentication with role-based access
- Centralised configuration management
- Auto-Discovery of server configurations 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" (mainframe term for a pool of like devices)
- 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

As an example, the following screen shot shows a configuration with the Luminex engines, virtual drive ranges, logical paths and tape pools, designed to give the storage administrator a clear, simple picture of all QuickRecover activity.

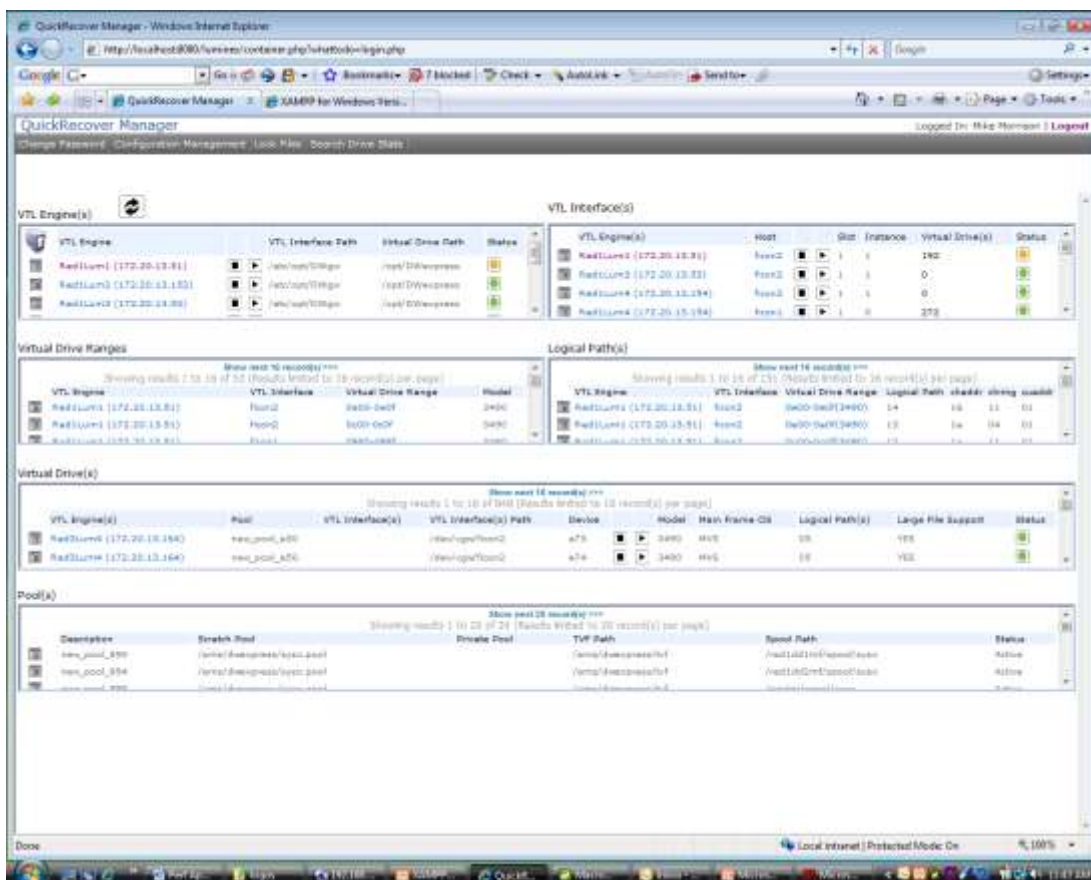


Figure 4- QuickRecover Management Software Web Interface



USER SUCCESS STORY

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 Level 2 data held offsite (an exposure in the event of primary site failure)
- 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 DISASTER RECOVER site
- 11:1 backup data deduplication achieved
- All former tape data now included in DISASTER RECOVER plan
- Much lower maintenance costs
- Redeployed tape staff
- Simplified and faster restores

SHODEN QUICKRECOVER PERFORMANCE:

Shoden's QuickRecover offers many benefits for mainframe environments, from cost-saving through to automated recovery, but there are significant opportunities to improve basic backup performance. The following three graphs (Figures 5, 6 & 7) show how QuickRecover's deduplication ratios improve over time, for different mainframe, UNIX and Wintel applications.

These results were derived from controlled environments; in each case the deduplication ratio (shown as the y axis) is plotted against the repeated backup runs. All 3 environments each have 3 applications (for example, CSB, CSC and CSC DB2 are mainframe applications). In practice with many additional repeated backups of the same data, we would expect to see the mainframe deduplication ratio climb towards 12:1, and for the open systems environments towards 20:1.




Figure 5- Deduplication ratio growth in three mainframe environments (UK FTSE100 Retailer)

Here three different mainframe applications (CSB, CSC, CSC DB2) were repeatedly backed up through the QuickRecover appliance. The first run will typically compress the data and will get about 4:1 compression; the second and subsequent runs will take advantage of Data Domains's deduplication techniques to compare blocks in flight; whenever an identical block is seen, a pointer is set, and the repeat block is discarded. Over time with repeated backups, the deduplication ratio in the mainframe environment will be expected to grow to between 10:1 and 12:1, typical figures seen in production with QuickRecover.

The following two graphs (Figures 6 & 7) show similar experiments for UNIX and Windows environments. Note the higher deduplication ratios for Windows and UNIX, even after only 3 or 4 repeated backups.




Figure 6- Deduplication ratio growth in three Windows environments (UK FTSE100 Retailer)

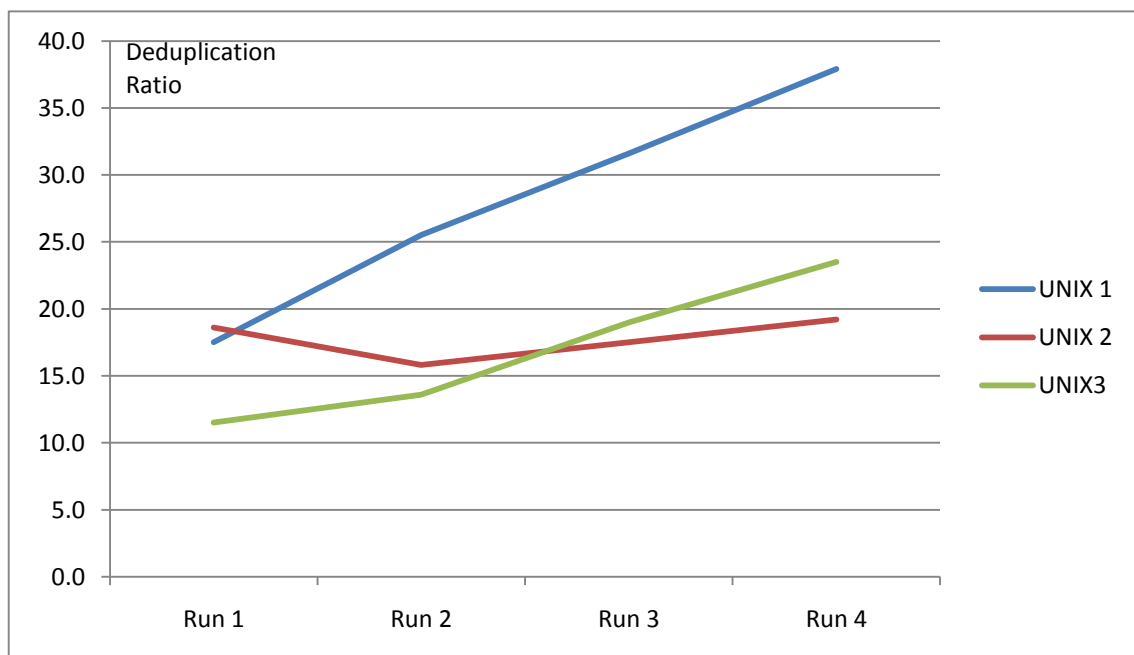


Figure 7- Deduplication ratio growth in different UNIX environments (UK FTSE100 Retailer)

Throughput improvement comparisons using QuickRecover:

As an example of the performance benefits, Shoden's QuickRecover was tested against a traditional tape environment in a large UK retailer's data centre, to gauge the advantage in backing up to disk (through deduplication) over tape. These were derived from a carefully controlled environment, with standard hardware and software being used. IBM's DFDSS backup utility was used in the first two tests, and in the third, CA's CA-Disk utility was used to migrate data directly to tape (9840) and Data Domain's disk storage (DD690).

- **DFDSS (single stream)**
 - Full volume dump (3390-3 = 2.7GB) to 9840 tape = **7.4 minutes**
 - Full volume dump (3390-3 = 2.7GB) to DD690 = **2.4 minutes**
- **DFDSS (15 parallel streams)**
 - Full volume dump (3390-3 = 2.7GB) to 9840 tape = **1 hr 25 minutes**
 - Full volume dump (3390-3 = 2.7GB) to DD690 = **31 minutes**

In the first test, a single 3390-3 volume was backed up to tape, and then to disk (behind Data Domain). In the second test, 15 3390-3 volumes were backed up in parallel. In both cases the elapsed times were reduced by more than 60% when the backup was switched from tape to disk.

- **CA-Disk Data Migration to Level 2**
 - Move 3.41GB to 9840 tape: = **5.0 minutes** (includes 27sec mount time)
 - Move 3.41GB to DD690: = **2.8 minutes**

In this test, CA Disk was used to migrate 3.41GB of user data from disk to tape (9840), and then the test repeated with the migrated data going instead to disk, behind the Data Domain appliance. Here, the elapsed time saving was more than 40%, after switching from tape to disk.

- **Replication simulation (asynchronous over Gigabit Ethernet)**
 - Input bandwidth = **250 MB/sec**
 - Output band-width (Post de-duplication) = **15MB/sec**

Data Domain's asynchronous replication minimises bandwidth use by only transferring changed and deduplicated data. In this environment, the bandwidth needed for replication is reduced by a factor of 16, while the elapsed times dropped by around 60% when switching from tape to disk.

BENEFITS OF THE SHODEN QUICKRECOVER™ APPROACH

Shoden's QuickRecover™ Mainframe VTL Appliance will provide the following benefits:

Reduced Total Cost of Ownership (TCO)

- Shoden QuickRecover eliminates tape media, minimises replication bandwidth and significantly decreases floor space, power and cooling requirements.
- Shoden QuickRecover removes the high maintenance costs associated with tape libraries.
- Shoden QuickRecover eradicates the management overheads involved in the administration, transportation, recycling and scratch pool management of physical tape.

Straightforward integration with existing backup environment

- The Shoden QuickRecover solution provides seamless integration to all current backup applications and does not require any changes or additional hardware or software.
- DataDomain is the only vendor to natively support the Open Storage option (OST) – This dramatically improves NetBackup 6.5 backup and restore performance.
- Support for all major Mainframe backup applications (CA1, DFSS, DFHSM, FDR etc.)
- Cost efficient - Data replication between primary and Disaster Recovery sites over existing IP network
- Data migration tasks from tape can be scheduled outside of backup and restore windows.
- Security - Provides a RAID protected backup and archive vehicle

Investment Protection

- Cost efficient - The Shoden QuickRecover (QR) backup solution is unique as it unifies both Mainframe and Open systems platforms with a common approach to backup
- Highly scalable data de-duplication solution designed to accommodate data growth
- Ability to replicate data asynchronously over any IP network.
- No legacy tape hardware is required for backward compatibility as all archive data is migrated to the new disk based solution.
- Early decommissioning of legacy tape infrastructure.

Environmentally friendly

- Dramatic floor space savings - 1 x 19" rack per site – allowing for growth and additional headroom for unanticipated requirements.
- Significant power, cooling and floor space savings allowing the tape robots and tape drives to be decommissioned immediately.
- Disk to disk remote replication significantly reducing bandwidth requirements.
- No need for "white van man" transportation of tapes to and from Disaster Recovery sites
- Time saving -Eliminates the process of staging tape based backups to disk and provides 'wire-speed' data de-duplication for backup and restore.

Shoden Value

- In depth experience of data migration for z/OS and Open systems platforms.
- QuickRecover management software simplifies backup and restore management whilst providing capacity planning, performance monitoring and failover automation in the event of component failure.
- Shoden is DataDomains' sole Mainframe technology partner
- Shoden has vast experience of delivering backup solutions and has full time in-house z/OS, Unix, Windows and technical resources.

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, and can easily be part of a unified backup mechanism spanning mainframe, UNIX and Wintel environments.

Philip Jones
Chief Technology Officer
Shoden Data Systems UK

Glossary

As with every aspect of IT, the mainframe has its own jargon and acronyms:

3390: Mainframe disk format used as basis for z/OS, introduced in 1989. 3390 has replaced the older 3380 format.

3480: Original mainframe tape format now superseded by 3490 and 3590. It was the first “square” or cartridge style tape to be used in the mainframe world.

3494: IBM tape library (high maintenance costs, power, heat etc)

3495: IBM tape library (high maintenance costs, power, heat etc); larger version of 3494.

Batch: Generic mainframe term for applications (“jobs”) that are not directly transaction based. Batch jobs are typically run overnight, where overall elapsed times are key to performance. The biggest batch job is usually backup, hence the importance of the backup “window”, the time during which the backups must run successfully, if they are not to overlap or interfere with online services.

DASD: Mainframe term for disk (Direct Access Storage Device)

Data set: Mainframe term for file, with a wide range of sequential and randomly accessed formats.

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. Deduplication is able to reduce the required storage capacity since only the unique data is stored.

DFDSS: IBM’s backup and recovery utility used for full volume and incremental backups.

DFHSM (usually HSM): IBM’s Hierarchical Storage Manager for archiving and migrating old, unused data. DFHSM usually has three levels, Primary or Level 0 (L0) on highest performing disk. L1 is usually an intermediate disk buffer, and L2 is usually on tape. Key to successful use of HSM is to get the correct size and the length of occupancy of the L1 pool, to minimise tape activity to and from L2.

ESCON: Original proprietary mainframe version of Fibre Channel, linking mainframe servers to peripherals. ESCON was launched by IBM in 1990.

FDR: Innovation Data Processing’s competitor to DFDSS

FICON: Proprietary mainframe version of Fibre Channel linking mainframe servers to peripherals. Previous version was ESCON

Full volume: the mainframe equivalent of a LUN. Traditionally all mainframe backups were based at the volume rather than file or data set level.

Incremental backup: Mainframe term for disk (Direct Access Storage Device)

MIPS: Millions of Instructions per Second. MIPS are a very old measure of processor performance, but still widely used. A typical mainframe environment today may have 2,000 – 20,000 MIPS.

MVS: Multiple Virtual Storages, the original name for the mainframe’s operating system, from 1974. MVS has now been superseded by z/OS, but all subsequent releases have built on the original MVS design.

RAID: Redundant Array of Independent Disks: a technology that allows computer users to achieve high levels of storage reliability from low-cost disk-drive components, via the technique of spreading data across multiple physical disks

Virtual tape library (VTL) is a data storage virtualisation technology used typically for backup and recovery purposes. A VTL presents a storage component (usually hard disk storage) as tape libraries or tape drives for use with existing backup software. Virtualising the disk storage as tape allows integration of VTLs with existing backup software and existing backup and recovery processes and policies. The benefits of such virtualisation include storage consolidation and faster data restore processes.

z/OS: Current mainframe operating system.

z-Series: Current mainframe hardware. Previous versions were S/360, S/370, S/390, 3090.