WHITEPAPER

DELL EMC DATA DOMAIN RETENTION LOCK SOFTWARE

A Detailed Review

ABSTRACT

Enterprises continue to see an exponential growth in the structured and unstructured data that is proliferating across their primary storage systems. Customers realize that the majority of this data is seldom accessed; yet they cannot delete this data given the compliance retention requirements for business records. As organizations drive formal adoption of archiving, IT administrators need cost-effective ways for their fast-growing archive storage needs, including compliance retention. This white paper introduces the Dell EMC Data Domain Retention Lock software that provides immutable file locking and data retention capabilities to meet a broad class of corporate governance and regulatory compliance (SEC 17a-4(f)) standards of archive data stored on Data Domain systems.

August, 2017

Table of Contents

EXECUTIVE SUMMARY 3
AUDIENCE 4
INTRODUCTION 4
SECURE RETENTION OF ARCHIVE DATA 4
GOVERNANCE ARCHIVE DATA REQUIREMENTS 4
COMPLIANCE ARCHIVE DATA REQUIREMENTS 5
TYPICAL DEPLOYMENT ENVIRONMENTS 5
DATA DOMAIN RETENTION LOCK SOFTWARE OVERVIEW 5
CONSOLIDATE GOVERNANCE AND COMPLIANCE ARCHIVE DATA 6
FILE LOCKING PROTOCOL 7
DATA DOMAIN RETENTION LOCK GOVERNANCE EDITION 8
SYSTEM MANAGEMENT 8
DATA DOMAIN RETENTION LOCK COMPLIANCE EDITION 9
“DUAL” SIGN-ON REQUIREMENTS 10
SECURE SYSTEM CLOCK 11
AUDIT LOGGING 12
LITIGATION HOLD 12
REGULATORY COMPLIANCE STANDARDS 12
TECHNICAL ASSESSMENT 13
SUPPORTED PROTOCOLS 13
CONSIDERATIONS FOR REPLICATING ARCHIVE DATA 13
DD RETENTION LOCK GOVERNANCE AND REPLICATION 14
DD RETENTION LOCK COMPLIANCE AND REPLICATION 15
DD RETENTION LOCK AND DD EXTENDED RETENTION 15
CONCLUSION 16

EXECUTIVE SUMMARY

Across the industry, enterprises continue to see an exponential growth in the structured and unstructured data that is proliferating across their primary storage systems. Customers realize that the majority of this data (as it ages with time) is not accessed often, yet they cannot delete this data because corporate governance and regulatory compliance (SEC17a-4(f)) standards mandate that data for business records must be securely retained for long periods of time (see Figure 1). As a result, companies are rapidly adopting formal archiving processes – so much so that the disk-based archiving market is forecasted to grow at a ~35% CAGR1 from 2010 through 2015.

Figure 1: Archive applications apply secure retention attributes

Almost all of enterprise data ranging from applications for processing content such as HR records or insurance document to traditional file/email records fall under strict retention guidelines. In addition, compliance retention policies continue to expand to include a broader variety of structured and unstructured data types. For optimal storage efficiency and data protection, customers require an archive storage system with:

- Support for both governance and compliance archive data with multiple retention periods on a single system
- Support for the majority of regulatory compliance standards including Security and Exchange Commission (SEC), Sarbanes-Oxley (SOX), Commodity Futures Trading Commission (CFTC), Food and Drug Administration (FDA), etc.
- Support for industry standard protocols (such as CIFS, NFS) for seamless integration with leading archive applications across various archive segments of file archive, email archive, enterprise content management (ECM) archive, database archive, etc.
- Next-generation protection storage that
    - Reduces the archive storage requirement with native compression and inline deduplication technology
    - Preserves all the data on the platform built of “storage of last resort” with the Data Domain Data Invulnerability Architecture
    - Enables consolidation of backup and archive data on a single system
    - Enables offsite protection via network-efficient replication

Dell EMC Data Domain Retention Lock® (DD Retention Lock) software provides immutable file locking and secure data retention capabilities for customers to meet both corporate governance and regulatory compliance standards, such as SEC 17a-4(f). DD Retention Lock provides the capability for IT administrators to apply retention policies at an individual file level. This software enables customers to leverage their existing Data Domain appliances to consolidate backup and archive data in accordance with governance and regulatory compliance standards.

[1] *IDC Report #230762, Archive Disk Based Storage Market: IDC WW Archival Storage Solutions 2010 – 2015 Forecast.*

AUDIENCE

This white paper is intended for Dell EMC customers, system engineers, partners, and members of the Dell EMC and partner professional services community who are interested in learning more about the DD Retention Lock software option.

INTRODUCTION

Data Domain Retention Lock software provides immutable file locking and secure data retention capabilities to meet a broad class of corporate governance and regulatory compliance (SEC 17a-4(f)) standards for archive data stored on Data Domain systems. This whitepaper explains the inner workings of both the Data Domain Retention Lock Governance edition and Data Domain Retention Lock Compliance edition.

This white paper will illustrate, how this software can be used to:

- Co-locate both governance and compliance data having different retention periods on the same Data Domain system
- Seamlessly integrate with your existing or new archiving application infrastructure to efficiently archive file/email, enterprise content management (ECM), database data and more
- Extend the use of your Data Domain system to consolidate archive data with backup data to maximize storage efficiency

This paper describes customer use cases, gives a product overview, and covers the interaction with other enterprise features such as replication. The security enhancements made specifically for the DD Retention Lock Compliance further allows deployment of strict compliant archive data along with governance archive and backup data on the same Data Domain system.

SECURE RETENTION OF ARCHIVE DATA

Unlike backup data, which is a secondary copy of data for recovery purposes, archive data is a primary copy of the data retained for long-term retention, secure and compliance retention purposes. As data ages and is seldom accessed, this data should be moved to archive storage, where it can still be accessed, but no longer occupies valuable primary storage space.

Archive data is usually stored for long-term with compliance retention policies and occasionally retrieved for eDiscovery needs. Since archive data is the primary copy of a data, IT administrators must ensure that the integrity of the data meets corporate governance rules and regulatory compliance (e.g. SEC 17a-4(f), etc.) standards.

GOVERNANCE ARCHIVE DATA REQUIREMENTS

Corporate governance standards for secure archive data retention are generally considered to be lenient in nature – allowing for flexible control of retention policies but not at the expense of integrity of the data during the retention period. These standards apply to environments where the system administrator is trusted with his or her administrative actions. The storage system has to securely retain archive data per corporate governance standards and needs to meet the following set of requirements:

- Allow archive files to be committed for a specific period of time during which the contents of the secured file cannot be deleted or modified
- Allow for deletion of the retained data once the retention period expires
- Allow for seamless integration with existing archiving application infrastructure through industry standard protocols such as CIFS and NFS
- Provide flexible retention policies such as allow extending the retention period of an archived file, revert of locked state of the archived file, etc.

- Ability to replicate both the retained archive files and retention period attribute to a destination site to meet the disaster recovery (DR) needs for archive data

COMPLIANCE ARCHIVE DATA REQUIREMENTS

The records retention requirements stipulated by the Securities & Exchange Commission ("SEC") Rule 17a-4(f) that defines compliance standards for archive storage expressly allows records to be retained on electronic storage media, subject to meeting certain conditions. Specifically, the conditions and requirements that an archive storage system must meet to be SEC compliant are:

- Preserve the records exclusively in a non-rewritable, non-erasable format. Specifically, as defined in the Rule itself, this requirement "is designed to ensure that electronic records are capable of being accurately reproduced for later reference by maintaining the records in unalterable form."
- Verify automatically the quality and accuracy of the storage media recording process
- Serialize the original, and if applicable, duplicate units of storage media, and the time-date for the required retention period for information placed on such electronic storage media
- Store separately from the original a duplicate copy of the record stored on any medium acceptable under 240.17a-4 for the time required

DD Retention Lock Compliance edition when deployed on the Dell EMC Data Domain storage system meets all of the above-mentioned SEC requirements set forth in Rule 17a-4(f), which expressly allows archive data to be retained on electronic storage media and meet strictest compliance requirements.

TYPICAL DEPLOYMENT ENVIRONMENTS

Enterprises have a slew of applications ranging from email servers to rich content management applications – and, the data across these application environments need to either meet governance or compliance retention. Below are some typical deployment environments:

- Archiving environments that are enforced for secure data retention requirements in line with corporate governance standards on existing or new archive data
- Archiving environments, specifically in industry verticals such as financial services, finance and banking, healthcare and pharmaceutical firms, legal and law firms that need to retain enterprise data in accordance with the strict retention requirements set forth by regulatory compliance standards (SEC 17a-4(f), etc.)
- Backup environments that are looking to store and securely retain archive data on the same deduplication based storage to drive further storage efficiency across the complementary workloads of backup and archive
- Environments where a net new end-to-end archiving solution is being architected and is required to seamlessly integrate with leading archive applications – such as Dell EMC SourceOne or Veritas Enterprise Vault
- Environments that want to reduce the primary storage footprint and spend by archiving inactive and aged data from primary storage system

DATA DOMAIN RETENTION LOCK SOFTWARE OVERVIEW

Dell EMC Data Domain Retention Lock software allows storage administrators, storage administrators, and compliance officers to meet data retention requirements for archive data when stored on a Data Domain system. DD Retention Lock software prevents files from being modified or deleted for a user-defined retention period.

Once the retention period expires, files can be deleted by the application, but cannot be modified. Files that are written to a Data Domain system but are not committed to be retained can be modified or deleted at any time. DD Retention Lock software comes in two editions – Dell EMC Data Domain Retention Lock Governance edition and Dell EMC Data Domain Retention Lock Compliance edition (see Figure 2).

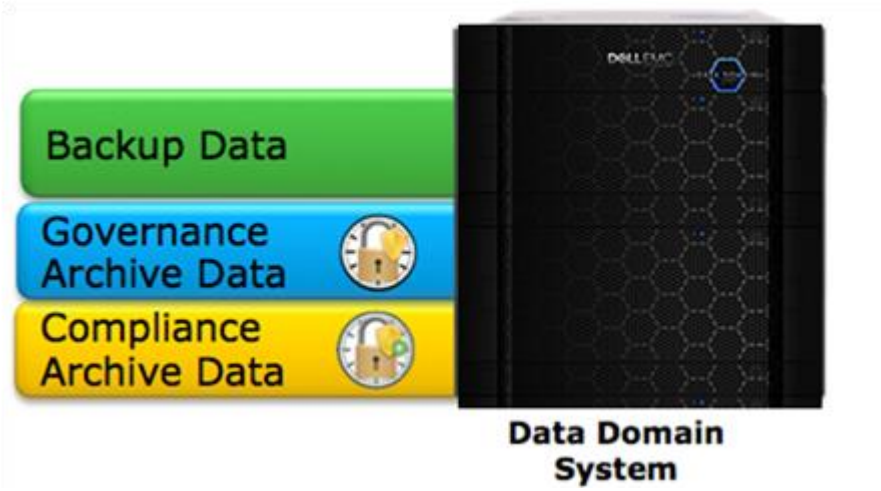

Figure 2: DD Retention Lock Governance edition and DD Retention Lock Compliance edition can *coexist* on the same Data Domain system.

- DD Retention Lock Governance edition maintains the integrity of the archive data with the assumption that the system administrator is generally trusted and thus any actions taken by the system administrator are valid to maintain the integrity of the archive data.
- DD Retention Lock Compliance edition is designed to meet strict regulatory compliance standards such as those of Security and Exchange Commission for 17a-4 Records (SEC 17a-4(f)).

CONSOLIDATE GOVERNANCE AND COMPLIANCE ARCHIVE DATA

Customers are looking for a storage solution that can consolidate both governance archive and compliance archive data on a single storage system. They want the storage system to be very easy to configure with granular management with built-in protection to help prevent common administrators mistakes.

To facilitate this consolidation of varied archive retention needs, a Data Domain Managed Tree (MTree) is used. MTree(s) are user-defined logical partitions of the Data Domain file system that enable granular management of data stored on a Data Domain system. Customers can enable DD Retention Lock software at an individual MTree level. A Data Domain system with DD Retention Lock Compliance can be configured to have one or more MTree(s) as compliance MTree(s) and/or with DD Retention Lock Governance can be configured to have one or more MTree(s) as governance MTree(s) (see Figure 3).

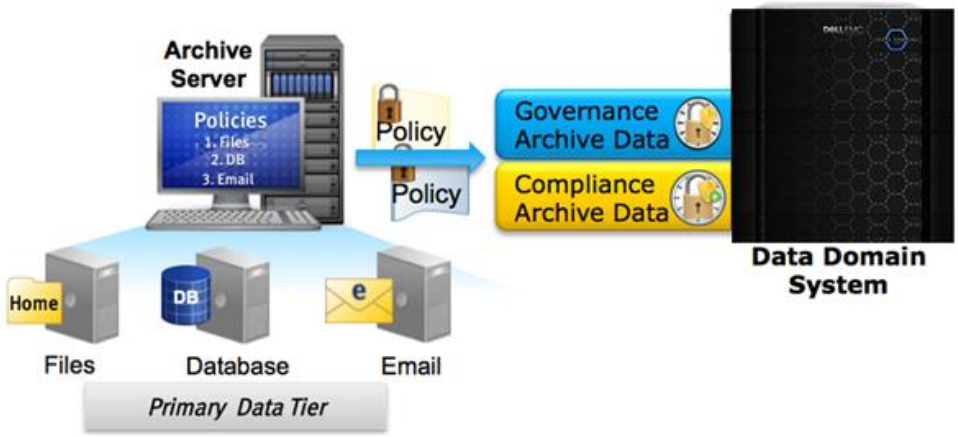

Figure 3: Data Domain Retention Lock software allows customers to apply different retention policies specific to the archive data type such as files, emails, database records, ECM data, etc.

This enables customers to be able to deploy both DD Retention Lock editions on the same Data Domain system[2].

This flexible deployment model allows customers to apply different retention periods for different types of archive data stored on MTrees (i.e. file, email, ECM, database archive, etc.) and meet both corporate governance and regulatory compliance standards on the same Data Domain system.

Customers must enable DD Retention Lock at an individual MTree level before any archive data stored on that MTree can be locked for governance or compliance retention. Before an MTree can be enabled with DD Retention Lock Compliance edition, the Data Domain system has to be configured for additional security measures (outline below). These measures ensure that administrative actions that could compromise the integrity of records are not under the control of just one administrative person. For specific instructions on configuring and enabling DD Retention Lock software, please refer to the *Dell EMC Data Domain Operating System Admin Guide*.

FILE LOCKING PROTOCOL

Once an archive file has been migrated onto a Data Domain system, it is the responsibility of the archiving application (or, manual scripts) to set and communicate the retention period attribute of the archive file to the Data Domain system. The archiving application sends the retention period attribute over standard industry protocols (CIFS, NFS), then the Data Domain system will enforce that retention period with DD Retention Lock.

The retention period attribute used by the archiving application is the last access time: the “atime”. DD Retention Lock software allows granular management of retention periods on a file-by-file basis. As part of the configuration and administrative setup process of the DD Retention Lock software, a minimum and maximum time-based retention period for each MTree (Managed Tree) is established. This ensures that the (atime) retention expiration date for an archive file is not set below the minimum or above the maximum retention period.

Let’s consider an example – an archiving application stores an archive file on the Data Domain system and sets the last access time (atime) of the file to the desired retention time, that is, a point in time in the *future* at which the file may be deleted. The retention period (atime) specified for a file in the MTree must be *equal to or greater than* the minimum retention period and equal to or less than the maximum retention period for that MTree.

If the retention period from the archiving application is:

- Less than the current date/time, or
- Less than the minimum retention period per MTree, or
- Greater than the maximum retention period per MTree

Then an error condition (permission denied error, referred to as EACCESS – a standard POSIX error) will be returned to the archiving application thus providing additional protection. The only exception here is in the scenario where the retention period is less than the current time plus 12 hours (tolerance window), and then the atime update will be ignored without an error and the file will not be locked for secure retention on the Data Domain system.

The archiving application must set the atime value and DD Retention Lock must enforce it to avoid any modification or deletion of files under retention of the file on the Data Domain system. For example, Veritas Enterprise Vault (EV), a file/email/SharePoint archive application, archives records for a user-specified amount of time. When EV retention is in effect, these documents cannot be modified or deleted on the Data Domain system. When that time *expires*, Enterprise Vault can be set to automatically dispose of those records.

For specific best practices to securely retain archive data via Veritas Enterprise Vault on the Dell EMC Data Domain system please refer to *Dell EMC Data Domain and Veritas Enterprise Vault Integration Guide*. For Dell EMC SourceOne, please refer to *Dell EMC Data Domain* and *Dell EMC SourceOne Integration Guide*.

[2] Note, that a single MTree cannot be configured as both governance and compliance at the same time.

DATA DOMAIN RETENTION LOCK GOVERNANCE EDITION

DD Retention Lock Governance edition allows customers to maintain the integrity of the archive data with the assumption that the system administrator is generally trusted with all legal actions performed on the Data Domain system (see Figure 4).

By enabling DD Retention Lock Governance edition on an MTree, IT administrators can:

- Apply retention policies at an individual file level of the data set on the Governance enabled MTree for a specific period of time
- Delete an archive file via an archiving application after the retention period expires
- Update the default values of minimum and maximum retention periods per MTree
- Extend the retention time of locked archive files

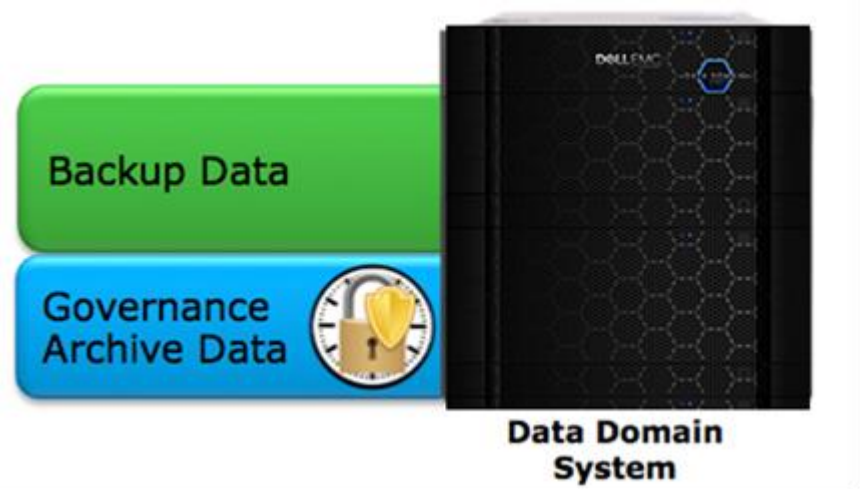

Figure 4: A Data Domain system can retain both Backup and Archive data that has to be retained per the corporate governance policies.

Locked files *cannot* be modified on the Data Domain system *even after* the retention period for the file expires. Archive data that is retained on the Data Domain system is **not** deleted automatically when the retention period expires; an archiving application must delete the file.

With DD Retention Lock Governance edition, IT administrators can meet secure data retention requirements while keeping the ability to update the retention period should the corporate governance policies change.

For example, an IT Administrator might want to:

- Revert the locked state of a file on a specified path name inside of an MTree
- Delete an MTree enabled with DD Retention Lock Governance

SYSTEM MANAGEMENT

Data Domain has designed an easy to use GUI in the Dell EMC Data Domain System Manager to help administrators monitor archived data. Alternatively, administrators can use available system commands (CLI).

Using the DD System Manager, customers can install the DD Retention Lock Governance license on the Data Domain system and can then enable DD Retention Lock Governance. DD System Manager provides the capability to update and modify the minimum and maximum retention period for MTrees. The DD System Manager GUI displays the various

states of the specific retention locked fields on the Data Domain system such as Unique Identification Number (UUID), etc. (see Figure 5 below)

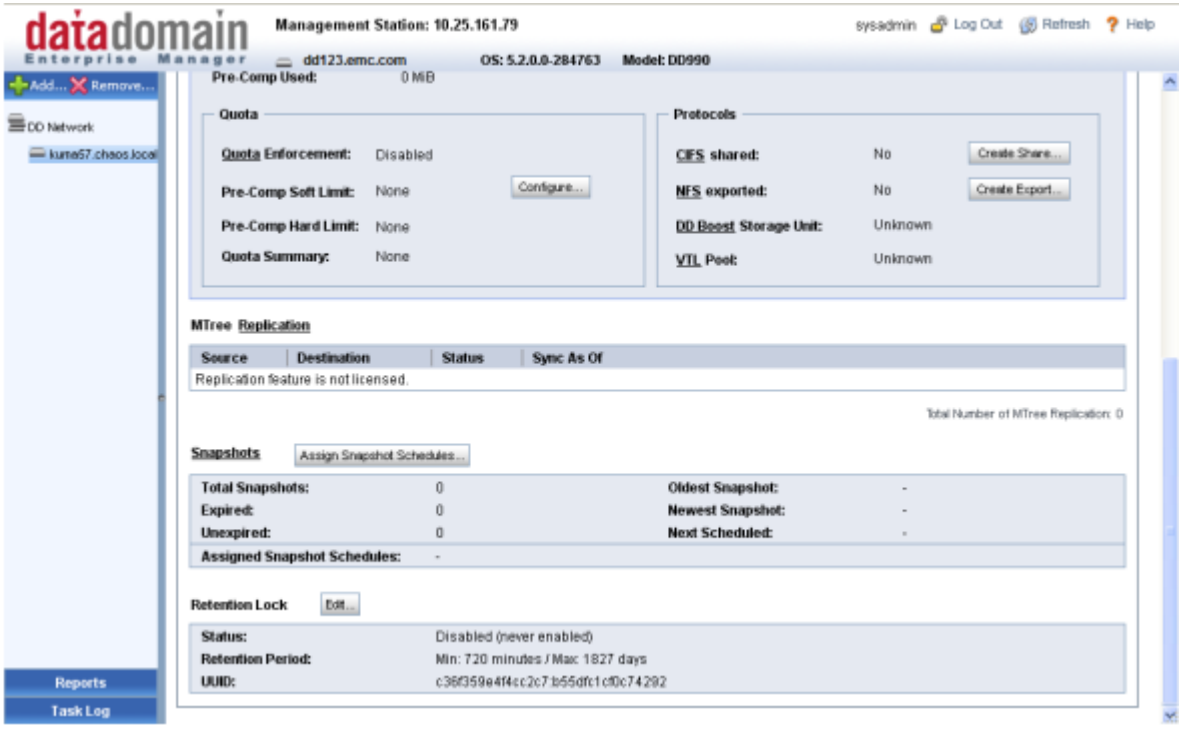

Figure 5: Data Domain System Manager showing Retention Period and UUID stat.

Please refer to the *Dell EMC Data Domain Operating System Admin Guide* and *Dell EMC Data Domain Operating System Command Reference Guide* for a detailed list and instructions on the full set of capabilities that are available for DD Retention Lock Governance edition on a Data Domain system.

DATA DOMAIN RETENTION LOCK COMPLIANCE EDITION

The DD Retention Lock Compliance edition meets the strict retention requirements of regulatory standards for electronic records such as SEC 17a-4(f) and other compliance standards that are practiced worldwide across industry verticals.

DD Retention Lock Compliance, when enabled on an MTree, ensures that all the files locked by an archiving application, for a time-based retention period cannot be deleted or overwritten under any circumstances until the retention period expires (see Figure 6). This is achieved via multiple hardening procedures such as:

- Requiring “dual” sign-on for certain administrative actions
- Completely disallowing operations that could compromise the state of locked and retained archive data
- Securing the system clock from illegal updates
- Audit logging for any operations that are executed upon the locked archive data
- Litigation hold allowing extension of the retention time of locked archive files
- Disabling various “doors” of access where someone could compromise the state of the locked data or the state of the retention attributes

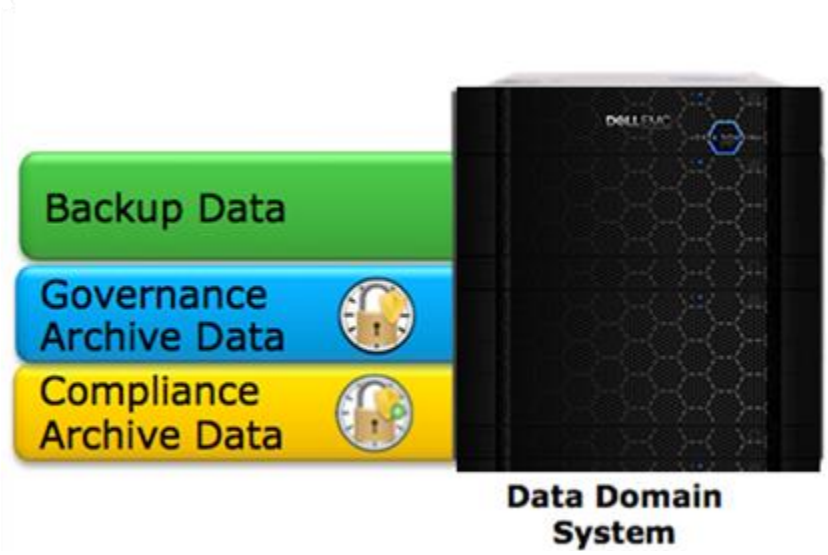

Figure 6: A Data Domain system that is being used for backup and governance archive data can be configured for the additional use case of archive data that needs to meet regulatory compliance (SEC 17a-4f) requirements.

"DUAL" SIGN-ON REQUIREMENTS

The most stringent requirement from compliance standards (such as SEC 17a-4(f)) is to ensure that any actions that could compromise the integrity of archive files prior to expiration of the retention period can only be executed by deliberate physical destruction methods.

To meet this requirement, DD Retention Lock Compliance edition provides a "dual" sign-on capability. This requires sign-on by the regular system administrator plus sign-on by a second authorized person (also referred to as "Security Officer") to perform certain administrative actions (see Figure 7). This ensures certain administrative actions are under the purview and control of higher authority above and beyond the system administrator.

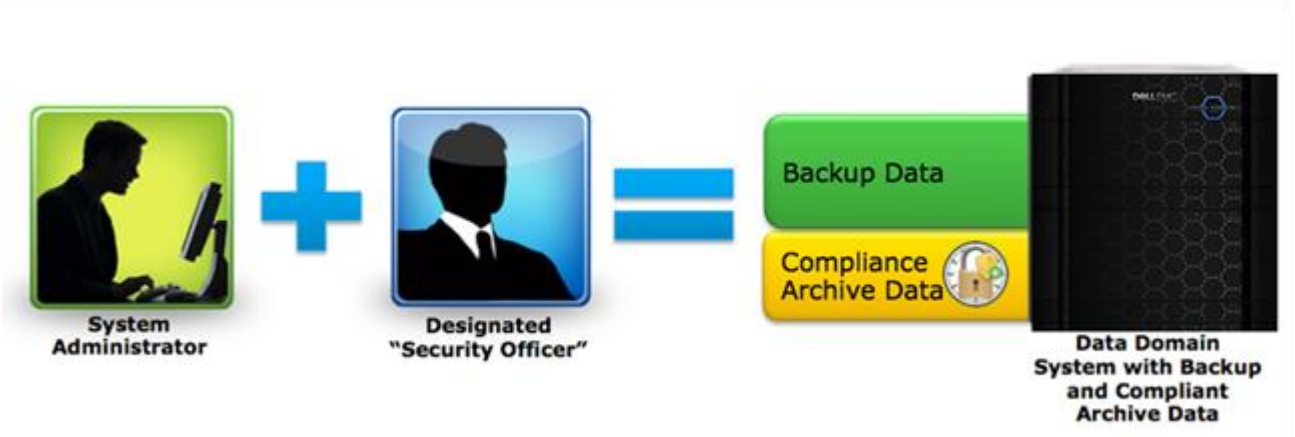

Figure 7: "Dual" sign-on capability of DD Retention Lock Compliance edition ensures that compliant archive data cannot be deleted under any circumstances.

It is possible to have *multiple* Security Officers configured on a single Data Domain system. Thereby, any one of the Security Officers can authorize system commands on the Data Domain system that require Security Officer credentials. It is important to note that the system administrator creates the first Security Officer; and after that only that designated Security Officer can add more Security Officers as authorized users on the Data Domain system. Having multiple

Security Officers is important for scenarios when one or more Security Officer is not available and certain critical operations have to be performed. In addition, creating multiple Security Officers in a single Data Domain system can prevent misuse of Security Officer authorization. Please refer to *Technical Assessment and Report from Cohasset Associates*, industry leading records management firm, on Security Officer model of DD Retention Lock Compliance edition.

Specifically, the primary administrative actions that require a dual sign-on are:

- Extending minimum or maximum retention periods of the MTree
- Renaming the MTree
- Deleting the installed DD Retention Lock Compliance license from the Data Domain system
- Other system support or maintenance actions that could potentially compromise the integrity of stored record files where the retention period has not expired

Additionally, the following are the operations or system commands that are completely locked down and therefore **cannot** be executed by anyone on the Data Domain system that has a DD Retention Lock Compliance installed and enabled:

- Destroying the entire Data Domain file system
- Deleting an MTree with DD Retention Lock Compliance enabled – even if the MTree is empty and has no locked files stored on it
- Disabling DD Retention Lock Compliance on an MTree after it's been enabled
- Reverting the retention state of the locked files on an MTree with DD Retention Lock Compliance enabled

SECURE SYSTEM CLOCK

Time-based retention is one of the primary requirements of regulatory compliance standards. In order to meet this requirement, a Data Domain system needs to ensure that any undesired changes to the system clock cannot be executed. Specifically, DD Retention Lock Compliance prevents users from changing the system clock via either system commands (CLIs) or by using the DD System Manager.

To ensure that the system clock cannot be modified, DD Retention Lock Compliance:

- Requires Security Officer approval for any system commands that could change the system clock
- Makes the system clock value persistent in the underlying storage by periodically writing it in the metadata file on the system
    - Then it will continuously check the current system clock time against this persistent system clock value
    - If the current time on the system clock is not within "acceptable" bounds (15 minutes) from this persistent time information, then this is considered as a *skew*
    - DD Retention Lock Compliance keeps track of the total skew for the current year and if the total skew becomes more than 2 weeks in that year, the system is locked down and the Data Domain file system will shut down
    - If the Data Domain system is locked down due to security clock violation, then it can only be resumed by providing Security Officer credentials

AUDIT LOGGING

Auditing capabilities are a requirement to meet compliance standards. Specifically, audit logs have to be kept for operations on locked data. This puts a requirement both on the archiving application and the Data Domain system. The logging of operations on the customer data is maintained by the archive application being used. Separately, the Data Domain system logs all management operations that affect locked files stored on the Data Domain system. All relevant operations are logged in a separate audit log that is available to the Security Officer. The system administrator cannot modify the audit log file on the Data Domain system that has DD Retention Lock Compliance edition installed and enabled.

LITIGATION HOLD

During periods of legal discovery, enterprises may be required by law to maintain their compliance data for extended periods of time. Dell EMC Data Domain systems provide this capability via litigation hold that allows the administrator to extend the retention lock periods on a per file basis beyond the maximum retention period of the containing MTree. Such an extension is allowed to a maximum period of 70 years from the current time. This capability can be driven via Archiving applications that support such extension capabilities.

REGULATORY COMPLIANCE STANDARDS

Compliance standards exist to verify that products comply with different regulatory standards across industry verticals. It's critical that customers ensure that a product used for secure retention of archive data receive a technical certification of standards compliance from a 3rd party with deep knowledge of regulatory standards and industry credibility.

In general, there are five United States federal regulations; the most notable being SEC Rule 17a-4(f). There is one international ISO standard and one European Union electronic records management guidance document. Refer to Table 1 below that lists various compliance regulations, industries impacted and the relevant DD Retention Lock edition that meets those requirements:

Compliance Regulation	Regulatory Agency	Industry/Vertical Impacted	Data Domain Retention Lock software
Sarbanes-Oxley (SOX)	Securities & Exchange Commission (SEC)	Public Companies	DD Retention Lock Compliance edition
SEC 17a-4(f)	Securities & Exchange Commission (SEC)	Financial Services	DD Retention Lock Compliance edition
21 CFR Part 11	Food and Drug Administration (FDA)	Pharmaceutical	DD Retention Lock software
CFTC Rule 1.31b	Commodity Futures Trading Commission	Financial Services	DD Retention Lock Compliance edition
HIPAA	US Health and Human Services	Healthcare Services	DD Retention Lock software
ISO Standard 15489-1	International Standards Organization	Public Companies	DD Retention Lock Compliance edition
MoREQ 2 (Model Requirements for the Management of Electronic Records)	European Commission	Public Companies	DD Retention Lock Compliance edition

Table 1: Summary of Regulatory Standards that DD Retention Lock software meets - from a Compliance Storage requirements perspective

TECHNICAL ASSESSMENT

Dell EMC engaged Cohasset Associates, an industry-leading records management consulting firm, for a independent and thorough technical assessment of the capabilities of the DD Retention Lock Compliance edition relative to meeting the strict requirements set forth in SEC Rule 17a-4(f) and a number of other regulatory compliance standards practiced worldwide as highlighted in Table 1 above.

Cohasset Associates performed an extensive technical due diligence on the features and functionality that are available via the DD Retention Lock Compliance software and certified that the Dell EMC Data Domain Retention Lock Compliance on the Data Domain system meets the relevant requirements of SEC 17a-4(f).

This means that during the SEC required retention period a Data Domain system with DD Retention Lock Compliance software:

- Provides the integrated control codes and record file management capabilities that ensures protection of record files from overwrite or erasure
- Provides for initial and ongoing accuracy and quality of the stored records
- Uniquely identifies each record file and duplicate copy
- Provides for a duplicate copy of the record files and recovery from the duplicate copy if required

In summary, the DD Retention Lock Compliance edition enables Data Domain systems to be the:

- Industry's *first* inline deduplication storage system that provides immutable file locking and secure data retention capabilities that meet a broad class of industry's strictest compliance standards for archive data
- Industry's *first* inline deduplication storage system enabling customers to deploy and co-locate both backup and archive data that has to meet compliance retention requirements

SUPPORTED PROTOCOLS

DD Retention Lock software is qualified and certified with industry leading archiving applications such as Dell EMC SourceOne, Dell EMC DiskXTender, Veritas Enterprise Vault (EV), etc. and is compatible via the industry-standard, NAS-based (CIFS, NFS) Write-Once-Read-Many (WORM) protocols. For a complete list of archiving and tiering applications that are qualified on Data Domain systems, please refer to the *Data Domain Archiving Applications Compatibility Matrix*.

Customers using backup applications such as Dell EMC NetWorker and Veritas NetBackup, can also use custom scripts to control the DD Retention Lock software on the Dell EMC Data Domain systems. For information on creating custom scripts to manage the retention policies of individual files, please see the *Dell EMC Data Domain Operating System Admin Guide* and refer to the section on "DD Retention Lock".

Note a Data Domain system with the Data Domain Virtual Tape Library software only supports DD Retention Lock Governance edition. Additionally, please note the following considerations for this configuration:

- Virtual tapes are represented as files on the Data Domain file system
- When customer creates a storage pool (a collection of tapes that map to a directory on the file system)
- Once created, one can use DD Replicator and DD Retention Lock on this MTree

CONSIDERATIONS FOR REPLICATING ARCHIVE DATA

Many companies have minimized the use of tape automation in their IT infrastructure by deploying deduplication storage for backup and operational recovery – Dell EMC Data Domain deduplication storage systems have been the market leaders in this category. In general, operational recovery includes retention periods from a few weeks to a few months. In addition, Data Domain systems also continue to revolutionize for backup and archive data that needs to be retained for longer period of times (years). By consolidating backup and archive data on a Data Domain system, storage

requirements can be reduced in size by 10 to 30x, making disk cost-effective for onsite retention, and highly efficient for network-based replication to disaster recovery sites.

Like most storage platforms, configuring disaster recovery is critical to a Data Domain system deployment. It is important to keep a full replica of all stored data in a separate system in a remote site that is protected from disasters and catastrophes. For Data Domain systems, Dell EMC Data Domain Replicator software provides simple, fast, robust WAN-based disaster recovery for the enterprise. It offers numerous replication types and policies and also supports a wide variety of topologies to meet the needs of various deployments.

DD RETENTION LOCK GOVERNANCE AND REPLICATION

For archive data that is locked for a specified period of time on a Data Domain system, it is critical for customers to be able to maintain the replicated copy of both the locked data and the retention attributes on the destination Data Domain system for DR scenarios.

Collection replication, MTree replication, and directory replication replicate (see Figure 8 below) the locked or unlocked state of files that are stored on the Governance enabled MTrees (see Table 2 below). This ensures that files that are locked on the source system remain locked after replication to the destination system. Only the source Data Domain system needs a DD Retention Lock Governance license for the locked data to be replicated and stored in the locked state on the destination Data Domain system.

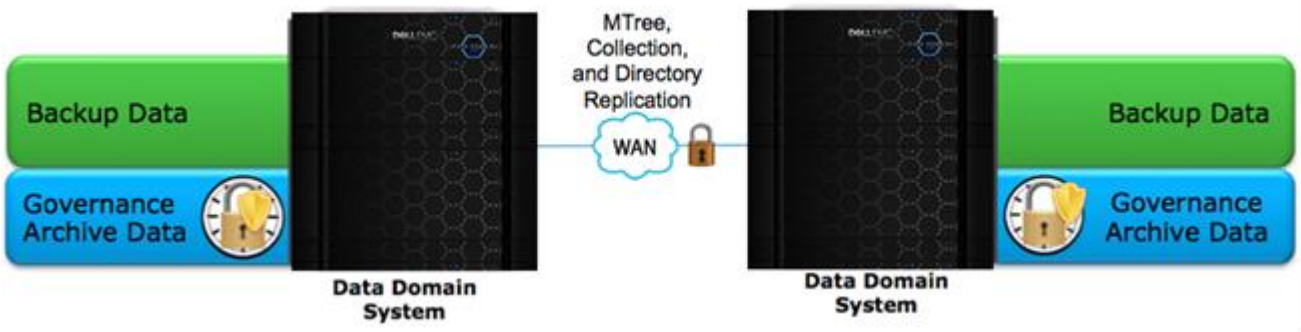

Figure 8: DD Replicator copies both the governance archive data under retention and associated retention periods from the source Data Domain system to the destination Data Domain system.

Replication Type	Support for replicating locked data	Replicate Min and Max Retenton Periods per MTree
Directory Replication	Yes	No
Collection Replication	Yes	Yes
MTree Replication	Yes	Yes

Table 2: Support for replicating locked data and metadata per Replication type

DD RETENTION LOCK COMPLIANCE AND REPLICATION

An absolute requirement for meeting compliance standards is to *"store separately from the original a duplicate copy of the record stored on any medium acceptable"*. Specifically, this rule from the SEC 17a-4(f) compliance standard requires that a separate copy of locked archive data must be stored on a secondary Data Domain system with the same retention attributes as the original. Therefore, DD Retention Lock Compliance must allow for replication of compliant archive data and make sure that both the source and destination systems meet compliance requirements (see Figure 9).

Figure 9: Either of MTree Replication or Collection Replication can be used to copy the backup and archive data from source Dell EMC Data Domain system to the destination Dell EMC Data Domain system

MTree replication or collection replication can be used to replicate the retention attributes of the locked archive files and associated MTrees to the destination Data Domain system. Here's an example scenario that expands on this capability for net new deployment of backup and archive (governance or compliance retention) data on a Dell EMC Data Domain system:

- Customer buys a pair of Data Domain systems and installs DD Replicator software, DD Retention Lock Governance software, and DD Retention Lock Compliance software.
- To consolidate both backup and archive data on the Data Domain system, the customer deploys backup data on MTree1, governance archive data on MTree2, and compliant archive data on MTree3.
- In this scenario, given that compliance archive data must be replicated to a destination Data Domain system, the customer uses MTree replication to meet disaster recovery requirements and consolidation needs of backup and archive data.

Furthermore, DD Retention Lock Compliance ensures that initialization, recover, and sync operations of collection replication will only proceed if the DD Retention Lock Compliance license is enabled on both the source and destination Data Domain systems. DD Retention Lock Compliance disallows replication from a compliant source Data Domain system to a destination Data Domain system that does not have DD Retention Lock Compliance enabled. Finally, DD Retention Lock Compliance ensures that any system commands or operations that either disable or break replication are under the purview of the Security Officer and can only be successfully executed after "dual" sign-on.

DD RETENTION LOCK AND DD EXTENDED RETENTION

DD Extended Retention software increases the storage scalability of a Data Domain system to enable cost-effective long-term retention of backup data on deduplicated disk. DD Retention Lock software is supported on DD systems that are enabled with DD Extended Retention software. Specifically, customers can install either DD Retention Lock Governance edition or DD Retention Lock Compliance edition or both on a Data Domain system with DD Extended Retention software.

On a Data Domain system with DD Extended Retention software installed, files that are locked on the active tier, with either DD Retention Lock Governance edition or DD Retention Lock Compliance edition, will remain locked when migrated to the retention tier. Once the retention period expires, the files on the retention tier can be deleted, but cannot be modified, and the associated space can be reclaimed starting with DD OS 5.3 release.

CONCLUSION

Customers continue to see exponential growth in structured and unstructured data that is proliferating across their primary storage systems. While the majority of this data is seldom accessed, it cannot be deleted due to governance or compliance (SEC 17a-4(f)) retention requirements. This has resulted in rapid adoption of formal archiving processes and impressive growth in the disk-based archive storage market. Therefore, many customers are looking to invest in deduplication based storage platforms to consolidate complementary workloads of backup and archive (governance and/or compliance) data to reap additional cost savings and storage efficiency.

Dell EMC® Data Domain® deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. By consolidating backup and archive data on a Data Domain system, storage requirements can be reduced in size by 10 to 30x, making disk cost-effective for onsite retention, and highly efficient for network-based replication to disaster recovery sites. Additionally, the system is protected by the Dell EMC Data Domain Data Invulnerability Architecture providing the industry's best defense against data integrity issues.

By deploying Dell EMC Data Domain Retention Lock® software on Data Domain systems, customers can securely manage the governance or compliance retention requirements. Data Domain Retention Lock software provides immutable file locking for both governance and compliance archive data sets, seamlessly integrates with leading archiving applications, and allows the consolidation of both governance and compliance archive data with different retention periods on the same Data Domain system. The DD Retention Lock Compliance capability in the Dell EMC Data Domain Retention Lock software enables Data Domain systems to be the industry's first inline deduplication storage system that meets a broad class of industry's strictest compliance standards for archive data.

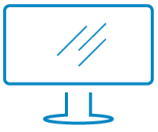

Learn more about Dell EMC Data Domain solutions

Contact a Dell EMC Expert