

IBM TSM DISASTER RECOVERY BEST PRACTICES WITH EMC DATA DOMAIN DEDUPLICATION STORAGE

Abstract

This white paper focuses on recovery of an IBM Tivoli Storage Manager (TSM) server and explores several architectural scenarios with EMC® Data Domain® Replicator software as a central mechanism for vaulting TSM to an alternate site for disaster recovery.

December 2010

Copyright © 2010 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Part Number h8128

Table of Contents

Executive summary.....	4
Introduction	4
Target audience.....	4
Data Domain's basic concept.....	4
Best practices recommendations	5
Data Domain replication considerations.....	6
Seeding the replication target system	6
TSM DR architecture overview	7
TSM database	7
TSM disaster recovery.....	8
TSM with Data Domain architectural scenarios	8
Data Domain system with physical tape vaulting (Creating Copy pool volumes)	8
Data Domain system with electronic tape vaulting.....	9
Two Data Domain systems with no physical tape library	11
Two TSM servers with bi-directional Data Domain replication	12
Multiple TSM servers with Data Domain replication	13
Conclusion.....	14
Appendix A: General sequence of TSM server recovery steps	14
TSM 5.5.x or earlier	14
TSM 6.0 or later	15
Appendix B: Process differences between physical and virtual tape libraries in TSM DR with Data Domain systems	16
Appendix C: Differences between tape and the FILE device class in TSM DR with Data Domain systems.....	17

Executive summary

EMC® Data Domain® deduplication storage systems provide unique and powerful disk-based backup and archive solutions with high-speed inline deduplication functionality for optimized backup and recovery design. In addition, the WAN-optimized replication of Data Domain systems provides effective offsite disaster recovery (DR) capabilities as a more reliable and cost-effective alternative to tape.

Introduction

The white paper provides technical and architectural information on the integration of IBM Tivoli Storage Manager (TSM) DR with Data Domain deduplication storage systems. It focuses on recovery of a TSM server and explores several architectural scenarios with EMC Data Domain Replicator software as a central mechanism for vaulting TSM to an alternate site for disaster recovery.

The reader should have a basic knowledge of TSM management practices as well as TSM Disaster Recovery Manager (DRM) procedures. A working knowledge of UNIX, Linux, and/or Microsoft Windows along with a basic understanding of the setup and management of Data Domain systems is also required. See the white paper [IBM TSM Backup with EMC Data Domain Deduplication Storage](#) for best practices in deploying Data Domain storage as a backup target in TSM environments.

The information in this paper can be applied to all supported versions of TSM. Variances to the recommendations for a specific version will be noted where applicable.

Target audience

EMC customers, system engineers, solution architects, and members of the EMC and partners' professional services community who are interested in DR configuration and best practices information when using Data Domain systems with TSM are encouraged to use this paper.

Data Domain's basic concept

Data Domain systems are typically used as a target for backup and archive data, due in part to the following capabilities:

- Support for most conventional backup applications through multiple protocols, for example, network-attached storage (NAS) for interfaces over Ethernet, and a VTL interface option over Fibre Channel.
- High-speed inline deduplication using small, variable-sized sequences to identify and eliminate redundant data before storing to disk.

- The EMC Data Domain Data Invulnerability Architecture provides integrated data protection technologies such as dual disk parity RAID 6, continuous recovery verification, with fault detection, and healing.
- Replication of deduplicated data over network-efficient wide area networks (WAN) to secondary systems for automated DR, which enables faster, “time-to-DR” readiness.

Data Domain deduplication storage systems come in a range of sizes and performance levels to support almost any TSM server environment. Usable capacity starts at just over 860 GB and scales to hundreds of terabytes. Logical capacities will vary due to variations in customer environments and the actual deduplication ratios achieved. Throughput rates can range from 450 GB/hr to over 8 TB/hr per appliance depending on the model and configuration. With data deduplication rates of 10x-30x on average, these systems are well suited to maintain multiple weeks of TSM backups from 17 TB, up to 7.1 PB of logical capacity in a single system.

Actual usage scenarios may vary so it is important to work with your Data Domain representatives to understand what results you can expect.

Best practices recommendations

The following are recommended best practices for TSM DR with Data Domain deduplication storage. These recommendations will be covered in detail throughout the remainder of this paper.

- Always perform TSM database backups as prescribed by best practices from IBM.
- Use the FILE device class for the TSM database backups for quicker, easier recoveries.
- When using the FILE device class for TSM database backups, use a dedicated device class using a unique directory on the Data Domain system.
- Use a dedicated storage pool for the TSM database backups.
- Do not use any form of native deduplication or compression within TSM for data sent to a Data Domain system.
- Do not configure any form of multiplexing.
- Use Tivoli Disaster Recovery Manager to assist with automation of TSM disaster recovery operations.
- Do not use any encryption before ingesting data into a Data Domain system. Data Domain has an option to encrypt data at rest on the system. Data in transit can be encrypted via the virtual private network, and other means.
- Dynamic tracking needs to be disabled when using the Data Domain system configured as a VTL.

- Set REUsedelay to the default settings of 0, and use a snapshot on the Data Domain replication target to replace the reusedelay functionality

Table 1. Recommended mount options for NFS

System environment	Settings*
Linux	intr,hard,rsize=32768,wsiz=32768,proto=tcp,vers=3,llock,combehind
AIX	intr,hard,rsize=32768,wsiz=32768,proto=tcp,vers=3,nolock
Other UNIX	intr,hard,rsize=32768,wsiz=32768,proto=tcp,vers=3,llock

*The recommended rsize/wsize is a minimum. Larger values may be used and have been seen to have beneficial results.

Data Domain replication considerations

Replication transmits data over Ethernet using TCP. The Data Domain system uses TCP port 2051 by default for replication.

The TCP port is configurable using the following CLI command:

```
# replication option set listen-port <VALUE>.
```

The Data Domain system does not reserve a physical Ethernet port for the replication feature; however, it is best practice to use the port with the IP address associated with the Data Domain system hostname.

Replication performance depends on daily change rates between new backup data and data residing on the Data Domain target system.

- The higher the change rate, the more data that needs to be moved within the replication window.
- Lower than expected deduplication rates will add time to replication and may exceed the replication window SLA.
- Network latency will have a significant impact on replication performance.
- WAN accelerators like the Cisco WAE help with packet loss when used in conjunction with DD Replicator. Any deduplication feature on a WAN accelerator should be disabled.

Seeding the replication target system

Initial replication (first data copy) is a time-consuming operation and will usually take significantly more time than later replications because the data is typically new and unknown to the target Data Domain system.

For initial replication, the goal is to facilitate the transfer of the large amount of unique data produced by initial backups ingested on the source shares/VTLs to the target Data Domain system. This is done by seeding the Data Domain target with the same data as on the Data Domain source so that only the minimum amount of data needs to be transferred between units on subsequent replications to maintain synchronization. The source and target systems should not be put into production until this seeding process is complete. Seeding is considered complete when the required replication window has been demonstrated to have been met. After seeding, the target unit can be shipped to the alternate location for replication.

TSM DR architecture overview

IBM Tivoli Storage Manager is based on a client/server architecture. TSM manages client backup data at the object level. The backup data originates from a client and may consist of a single file, directory, file system, database data via TDP's (Tivoli Data Protection) Application Agent, or partition images. Data is tracked at the object level and is stored in storage pools, which are groups of volumes of the same DEVICE type. TSM domains are used to group clients with policies that control retention and storage preferences for the client data. There are many variations of the basic configuration within the TSM Data Protection family of products. For more info with TSM concepts and terminology, please see the following documents:

- [IBM TSM Backup with EMC Data Domain Deduplication Storage](#)
- [IBM Tivoli Storage Manager, Version 6.2 Server Overview](#)

TSM database

With the release of TSM 6.0, TSM is now using DB2 to store information about registered client nodes, schedules, and the client data in storage pools. Before TSM 6.0, TSM was using an internal, proprietary database. The TSM server keeps this database protected through transaction logs and periodic backups. A TSM database backup can be full, incremental, or a snapshot, and captures all transactions currently committed to the database. In any TSM iteration, the TSM database is key to the operation of the TSM server. The TSM database cannot be recovered from the backed-up client objects. Thus it is important to always do TSM database backups according to IBM best practices. Typically, database backups are done at least twice each day. The full and incrementals are held onsite for operational recovery and the snapshot is sent offsite for DR purposes.

One drawback of TSM database backup to tape is that each backup must be done on a dedicated tape media. This leads to inefficient use of tape capacity and adds to the management of tape media. In addition to tracking tape media for client/app backups, administrators must also track the database backup media. With multiple database backups per day, and maintaining only 5 days of database retention, each TSM instance will require a minimum set of 10 significantly underutilized tapes.

With the introduction of Data Domain deduplication storage to the TSM operational schedules, TSM database backups can be written directly to the Data Domain system using media in a FILE or LTO device class volume. With Data Domain deduplication technology, administrators avoid the problem of underutilized capacity with tape. After the first database backup, only new, unique data is stored on the Data Domain system. Data Domain storage does not support putting the TSM database/logs directly on the Data Domain system. Use Data Domain deduplication storage for the TSM database backup and TSM data storage. The TSM database/logs should live on primary storage (designed for random access).

EMC recommends the use of a FILE device class for database backups on Data Domain systems for customer environments that support such configurations (that is, Ethernet connectivity to the Data Domain system). VTL can also be used to back up the TSM database. TSM database recovery from FILE device class volumes are much easier and convenient. Less time is spent configuring and reconfiguring an environment in preparation for restore via a FILE device class than traditional library/VTL recovery methods.

If a second Data Domain system is introduced at an alternate site, TSM database backups can be replicated to it. DD Replicator eliminates the need for TSM to create a DB snapshot solely for offsite storage and recovery. The offsite recovery can occur from the TSM database full backups. Because of the elimination of the database snapshot, more frequent TSM database fulls could potentially be obtained for more frequent recovery points. As long as all Primary pool data is stored on the Data Domain system, TSM Copy pools are also no longer needed. Note that TSM is not aware of the replicated data. All customer procedures for disaster recovery and data handling need to be reviewed in order to handle the differences between local and remote TSM environments.

TSM disaster recovery

There are several architectural choices that can be explored when introducing the Data Domain system for vaulting TSM data to an alternate site to ensure rapid and efficient disaster recovery.

It is always a best practice to make sure the TSM database, configuration files, and storage pools are available for TSM recovery operations.

TSM with Data Domain architectural scenarios

Data Domain system with physical tape vaulting (Creating Copy pool volumes)

In Figure 1, a Data Domain system is used as a primary storage pool replacing the traditional TSM DISK pool.

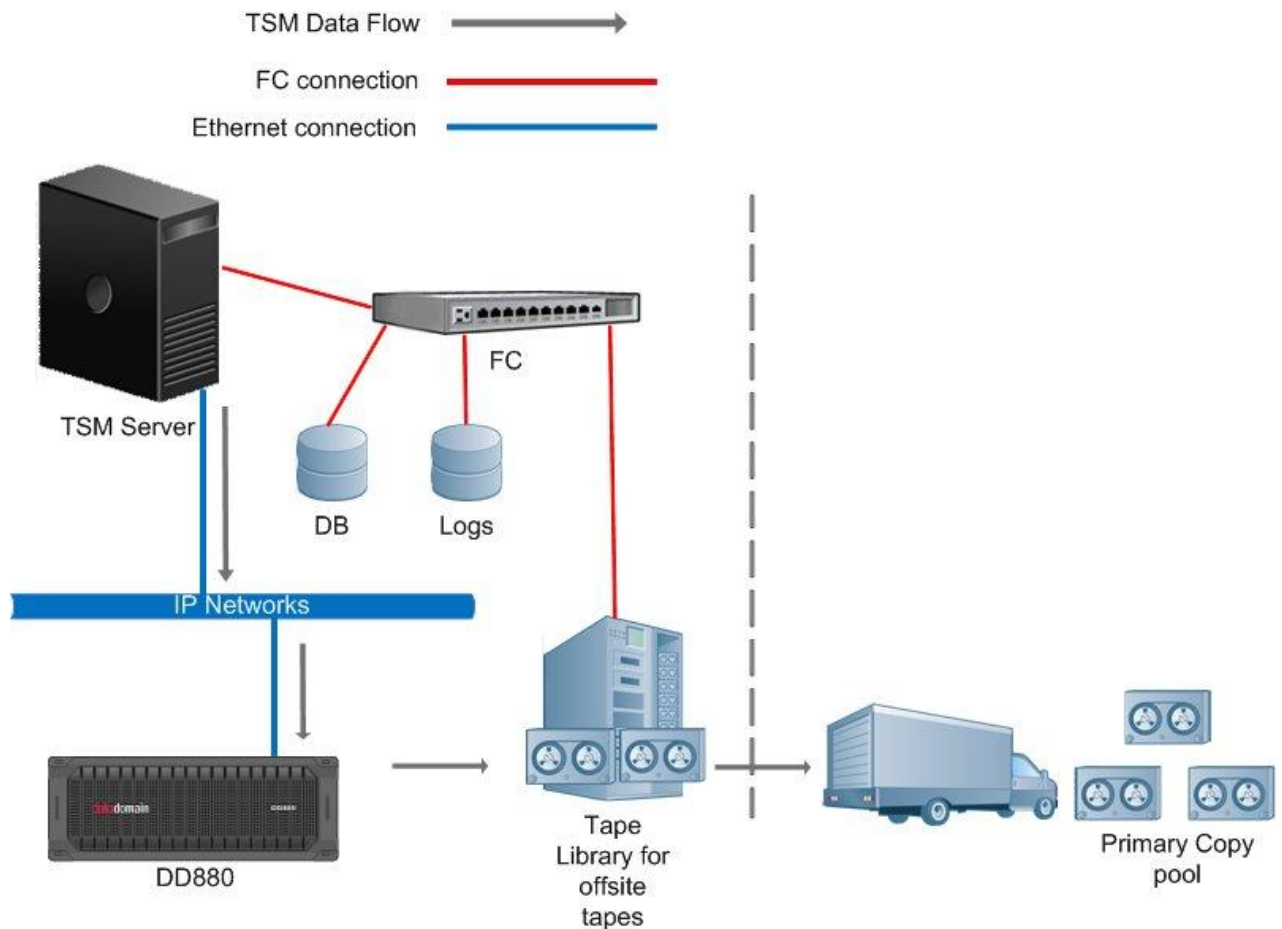


Figure 1. Local Data Domain system as a primary storage pool for local vaulting

A tape library is used as a Copy pool to create a secondary copy of the primary storage pool and the tapes can be manually shipped offsite. In this scenario, the migration step has been eliminated when a Data Domain system is used as the primary storage pool, shortening the daily operational workload.

Note: When the TSM data is copied from the Data Domain system to a tape library Copy pool, the data sent to the physical library is not deduplicated or compressed; the TSM data is transmitted in a TSM native format.

Data Domain system with electronic tape vaulting

In the next figure, a Data Domain system is used as the primary storage, replacing TSM DISK or the tape primary storage pool. A tape library is located in the alternate site, which is used for electronic vaulting by TSM.

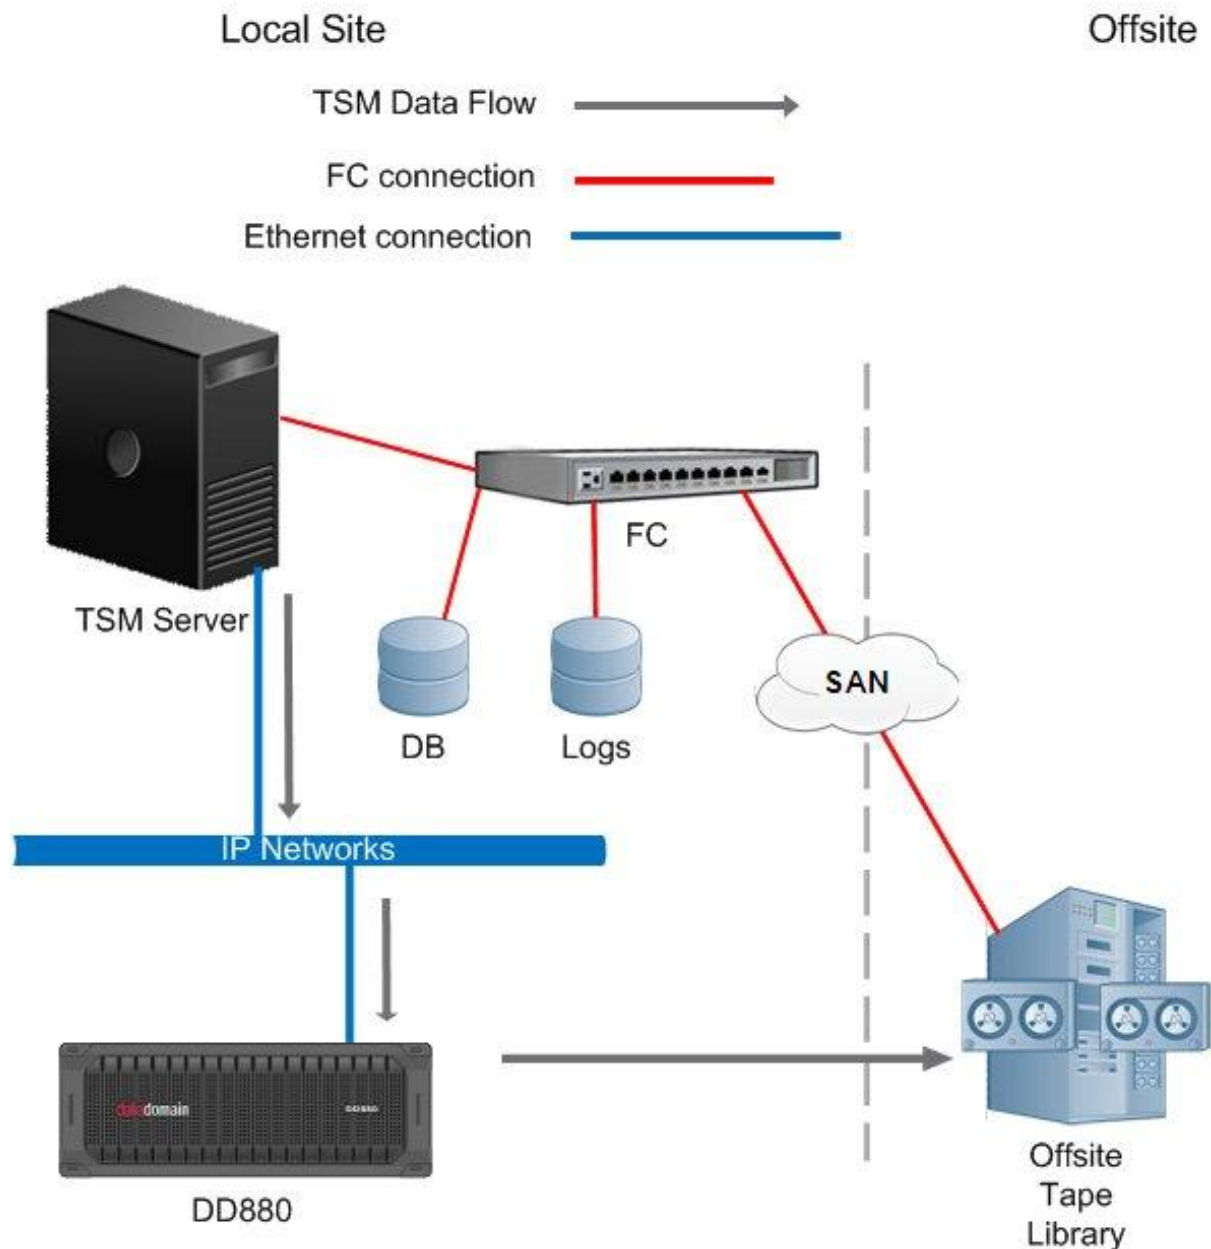


Figure 2. Local Data Domain system as a primary storage pool with TSM electronic vaulting

The migration requirement again has been removed from daily TSM activity, shortening the TSM daily operational workload. EMC recommends replacing the offsite tape library with a secondary Data Domain system as a replication target. With DD Replicator, there is no need to use an expensive SAN environment to have a secondary copy of TSM database and TSM data backups. The next section has more information.

Note: When the TSM data is copied from the Data Domain system to a tape library Copy pool, the data sent to the physical library is not deduplicated or compressed; the TSM data is transmitted in a TSM native format.

Two Data Domain systems with no physical tape library

In the next figure, the TSM offsite tape library is eliminated by using two Data Domain systems with DD Replicator. With the second Data Domain system located in an alternate site, the TSM database and data backups are replicated to the secondary Data Domain system as second copies. As the TSM database and data backups are processed on the primary Data Domain system, the unique segments and metadata representing each file in the TSM backups are replicated to the secondary Data Domain system, allowing the overall “time-to-DR” to be minimized. In many cases, replication is completed soon after the initial backup completes. By utilizing Data Domain replication, TSM daily operational tasks are reduced. Tape vaulting and migration are eliminated, and TSM database size is reduced.

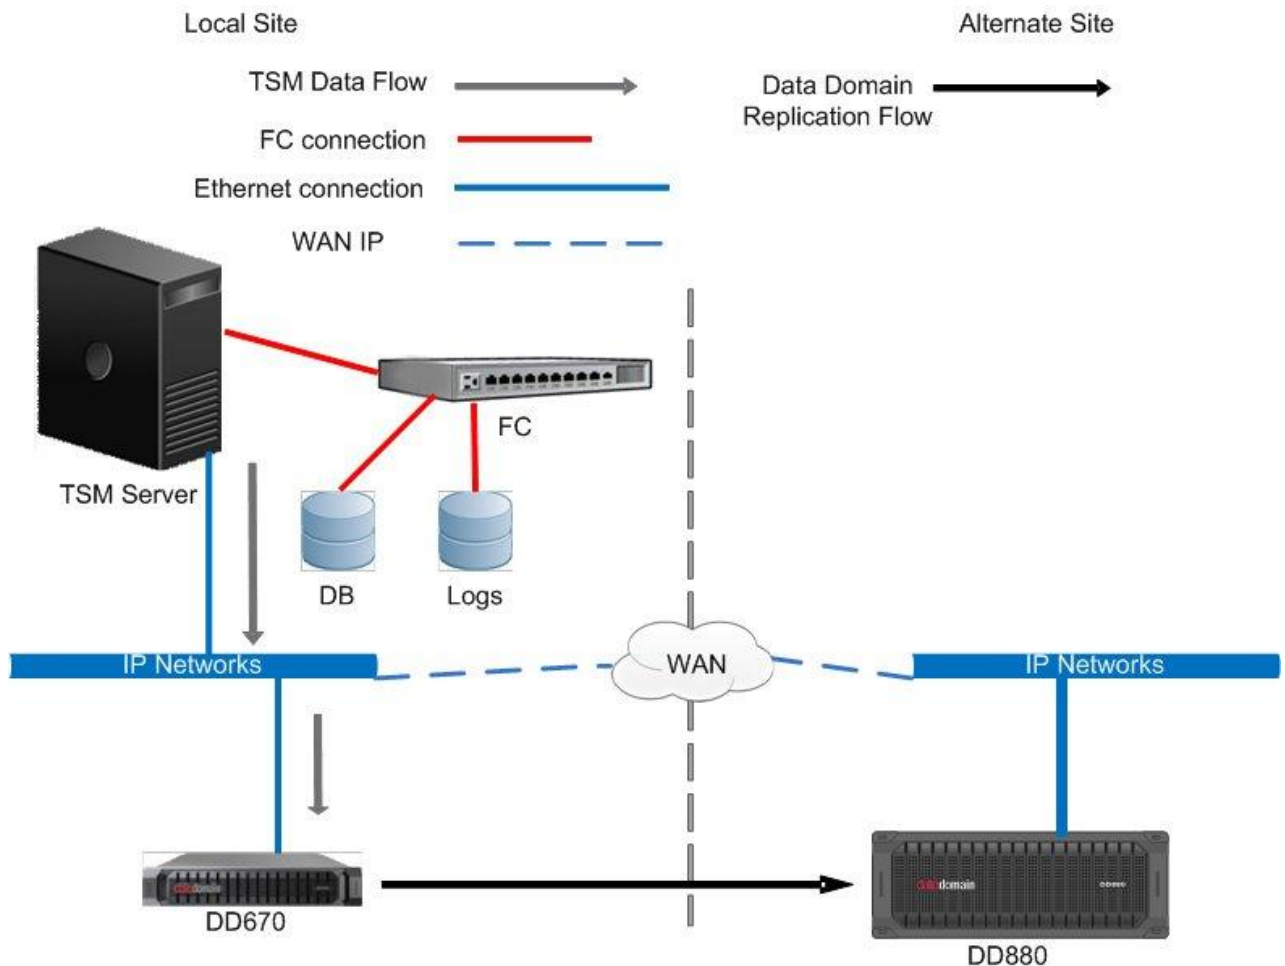


Figure 3. Data Domain replication to an alternate site

NOTE: The TSM server is not aware of the secondary copy because the duplicate copy was done through DD Replicator, which has no mechanism to update the TSM server database regarding the existence of the second copy.

The administrator will need to have the exact same server environment (operating system, TSM versions, and so on) at the disaster site and ensure that the replication

is in sync, including the database. The database will need to be restored first. All the devices need to be mounted and visible to the DR TSM server. For more details, reference [Appendix A](#).

Two TSM servers with bi-directional Data Domain replication

In the next figure, Data Domain systems support multiple TSM production servers at two different sites.

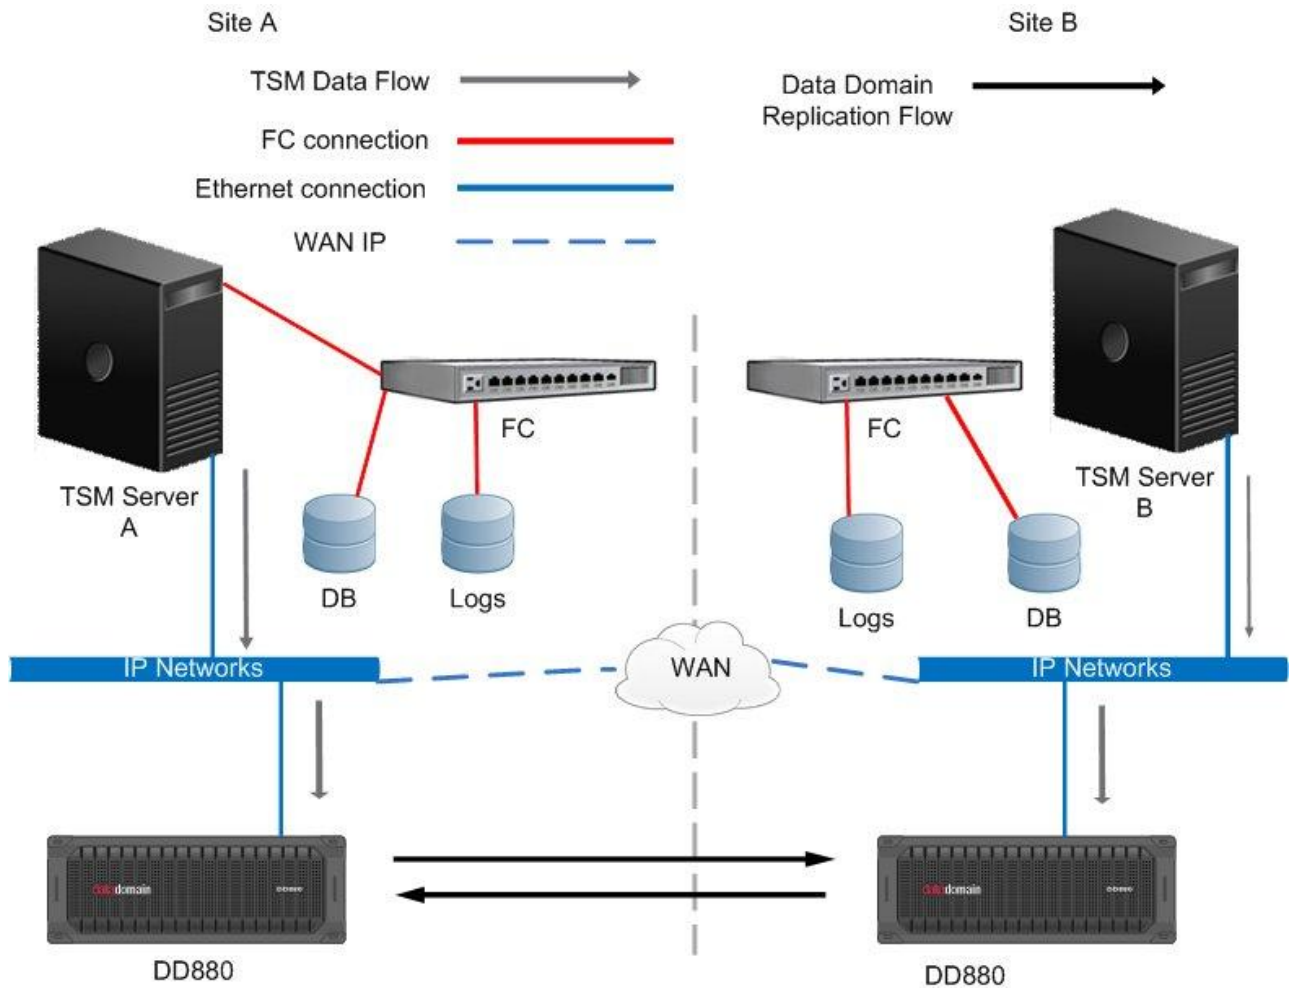


Figure 4. TSM with Data Domain bi-directional¹ replication

Each site has its own TSM server, and each TSM server has replicated copies of the TSM database backups and backup data coming from Data Domain replication. DD Replicator enables disaster recovery of each TSM server at each site because of the secondary copies of the TSM database backups and backup data.

NOTE: With bi-directional replication, the TSM servers are not aware of the secondary copies of each other's data because the duplicate copies were done through Data Domain Replicator software, which has no mechanism to update the TSM server

¹ Bi-directional replication will make copies from both Data Domain systems to each other.

databases regarding existence of the second copies. In a DR scenario, a recovery TSM server at the alternate site would have to be started up with the replicated copies of the TSM database backup and backup data residing in the alternate site Data Domain system.

Multiple TSM servers with Data Domain replication

In the next figure, a centralized disaster recovery site supports multiple TSM servers spread across different sites, by replicating all the TSM database backups and backup data to the centralized disaster recovery site using DD Replicator. Data Domain systems enable recovery of any of the TSM servers quickly at the central site.

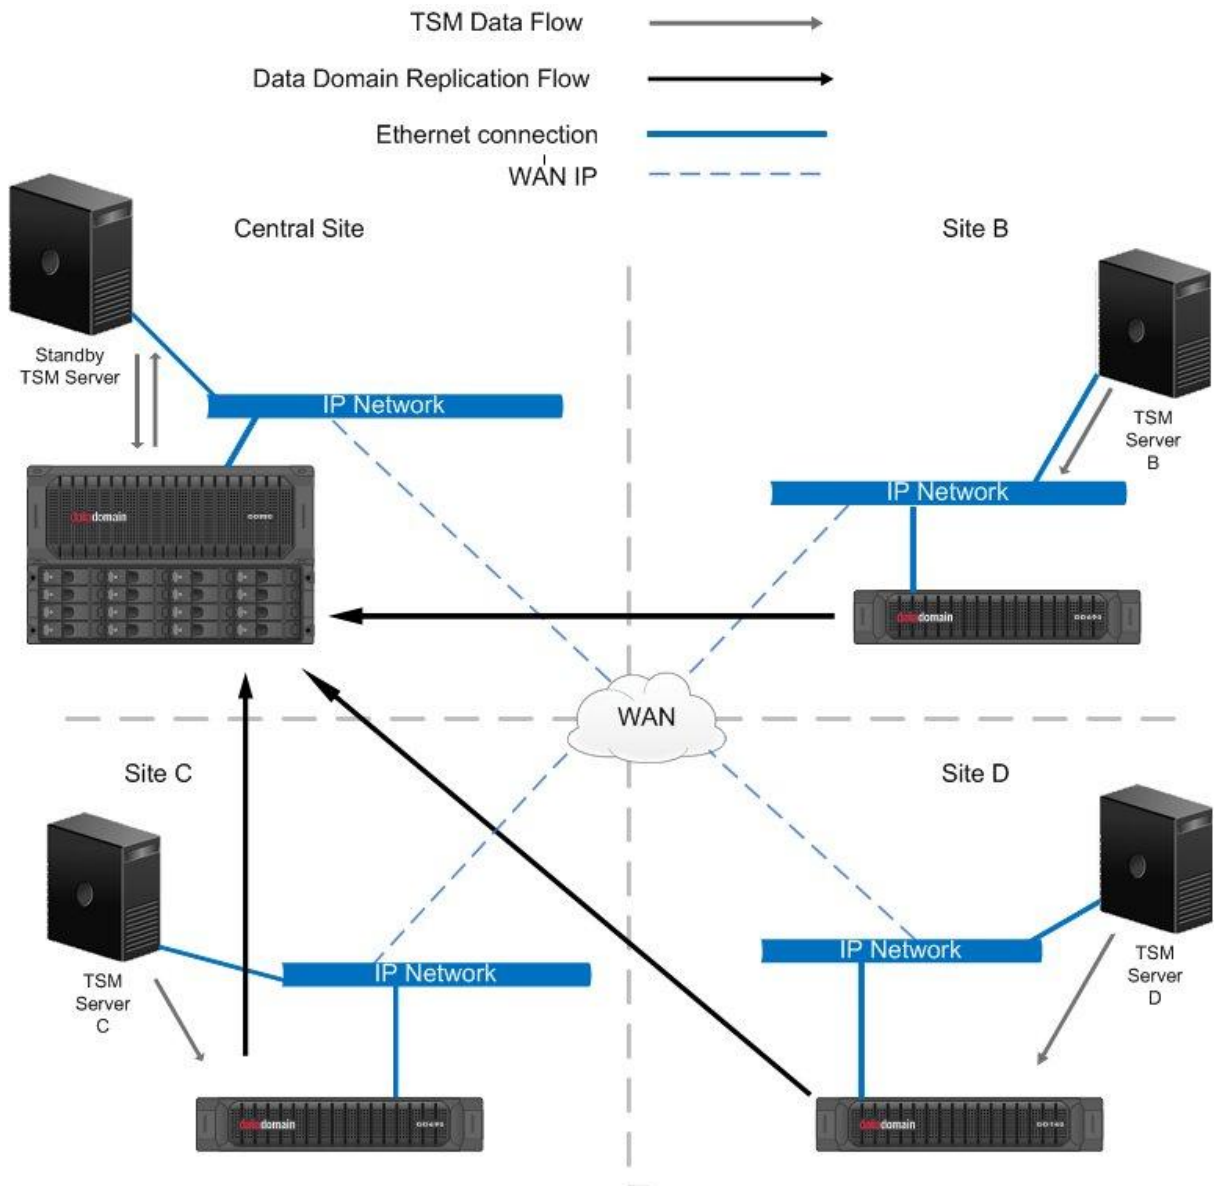


Figure 5. Multiple TSM servers with Data Domain replication

NOTE: As in the previous example (Figure 4 on page 12) the TSM servers depicted in Figure 5 are not aware of the secondary copies of each other's data because the duplicate copies were done through Data Domain Replicator software, which has no mechanism to update the TSM server databases regarding existence of the second copies. In a DR scenario, a recovery TSM server at the alternate site would have to be started up with the replicated copies of the TSM database backup and backup data residing in the alternate site Data Domain system.

Conclusion

Data Domain systems offer a variety of new storage architecture options for TSM configurations. Compared to a physical tape library, Data Domain deduplication storage removes the complexity and management overhead associated with tapes. With Data Domain replication, TSM disaster recovery which required significant architectural measures and management overhead is now easier to architect and manage. Data Domain replication provides a viable alternate strategy with TSM disaster recovery that couples a reduction in TSM daily administrative tasks.

Appendix A: General sequence of TSM server recovery steps

What follows is a list of steps that can be used to recover a TSM server instance in conjunction with Data Domain replication. EMC recommends using TSM DRM to automate TSM disaster recovery operations.

TSM 5.5.x or earlier

Assumptions: The standby server is running the same OS and patch level as the production TSM server and it can access the Data Domain replication target at the disaster recovery site.

1. Install the appropriate TSM version.
2. Ensure that the replication sync point between the Data Domain systems are achieved with each maintenance cycle after the TSM database backup and appropriate plan and configuration files have been placed in the replication target.
3. Once the replication sync point has been obtained, create a snapshot on the Data Domain system replication target.
4. At the time of the disaster recovery, on the Data Domain system replication target, fastcopy the appropriate directories or VTL pools from the appropriate snapshot.
5. Use the fastcopy command to copy the data from the Data Domain system replication target as the source for the recovery. (NOTE: Update the TSM recovery script to remove the steps that destroy/mark unavailable primary volumes.)

6. Configure the TSM server to access the Data Domain system using the NFS/CIFS/VTL protocols. The device class must be configured identical to the production system. See Table 1 on page 6 for the NFS mount option settings.
7. Define the TSM database and log file systems to match the exact configuration of the production TSM server.
8. Format the TSM database and log volumes.
9. Use “dsmserve restore db” to restore the TSM database.
10. Register the TSM server license.
11. Inventory required media for storage pools (this is only required if VTL or tape is managed by TSM).
12. Run a restore in the alternate backup environment to verify the recovered TSM server instance.

TSM 6.0 or later

Assumptions: The standby server is running the same OS and patch level as the production TSM server and it can access the Data Domain replication target at the disaster recovery site. The Data Domain system will need the appropriate TSM server permissions granted to the system /backup export/share. Required permission access is also required for the VTL device path on the OS, and the user who owns the TSM server instance must have the appropriate permission.

1. Install the appropriate TSM version.
2. Ensure that the replication sync point between the Data Domain systems is achieved with each maintenance cycle after the TSM database backup and appropriate plan and configuration files have been placed in the replication target.
3. Once the replication sync point has been obtained, create a snapshot on the Data Domain system replication target.
4. At the time of the disaster recovery, on the Data Domain system replication target, fastcopy the appropriate directories or VTL pools from the appropriate snapshot.
5. Use the fastcopied data from the Data Domain system replication target as the source for the recovery. (NOTE: Update the TSM recovery script to remove the steps that destroy/mark unavailable primary volumes.)
6. Create the database instance user ID and group as in the production server.
7. Create the database directories, active directories, and archive directories as in the production server.
8. Run the dsmicfgx utility to configure the replacement instance. This step configures the API for the DSMSEV RESTORE DB utility.
 - a. Specify the instance userid and password.

- b. Specify the database directories, active directories, and archive directories.
9. Remove the database instance that was created by the dsmsicfgx utility. For example:

```
# dsmserv removedb TSMDBA
```
10. Restore the original dsmserv.opt, volume history, and device configuration files to the instance directory.
11. Run the DMSERV RESTORE DB.

Appendix B: Process differences between physical and virtual tape libraries in TSM DR with Data Domain systems

Standard IBM TSM best practice for recovery should still be followed. When recovering TSM when using Data Domain systems in VTL mode, you must prepare volumes for use in the new VTL configuration and secure a positive sync point and static copy of the replicated data.

Once Data Domain replication is configured and daily operations are underway (see Chapter 10 in the [EMC Data Domain Administration Guide](#) for further information on Data Domain replication setup and configuration), the replication should be synchronized after the TSM database backup is taken and appropriate TSM configuration files are placed in the Data Domain replication contexts that will be used for the recovery. Furthermore, a snapshot on the Data Domain replication target system should be taken.

NOTE: In a VTL environment, VTL configuration is not replicated and must be configured on the Data Domain replication target, and need only be configured once, as the configuration is unaffected by future replication.

During the DR recovery process, a fastcopy of the replicated TSM data needs to be made. Fastcopy is a Data Domain OS command that allows making a secondary copy of data on the same system. For example, if the vault name in the Data Domain target replication were TSM_DB1_R, the fastcopy would be made to TSM_DB1. The volumes that exist in TSM_DB1 would then be imported to the secondary VTL on the Data Domain replication target system. After this, all of the typical TSM considerations of physical tape libraries apply – the device configuration file for the instance where TSM is being recovered must be modified to reflect the correct path to the library changer and at least one tape drive; the slot number of the database backup/snapshot volume must be updated in the volume history file, or else the volume list would need to be manually mounted in a tape drive via the Data Domain VTL GUI.

Once the TSM instance is recovered as with any TSM DR, all paths and definitions should be updated/refined to ensure that the TSM instance is addressing the library and drives appropriately.

Appendix C: Differences between tape and the FILE device class in TSM DR with Data Domain systems

The process to recover a TSM instance when using the FILE device class as opposed to using tape, whether virtual or physical, removes the requirement of configuring or reconfiguring a library and tape drives. The recovery process can use a direct physical path to the FILE device class volume whether the path to the volume at recovery matches production or not. Once recovery of the TSM database has been completed, as with any TSM DR, all paths and definitions should be updated to ensure that the TSM instance is addressing the device paths appropriately.