

WHITE PAPER

Terremark Backup and Restore Backup Lifecycle Management

EXECUTIVE SUMMARY

Within most corporate environments, data can be described as critical (i.e. integral to ongoing operations) or important (i.e. valuable but not mission critical). An example of the former is your email information store and an example of the latter is your financial data from five years ago. The Terremark Backup and Restore service provides an online, off-site solution that allows for fast retrieval of critical data. For important data, the Terremark Backup and Restore service provides Backup Lifecycle Management (BLM), a more-cost effective secondary storage solution that increases the time-to-restore but maintains the same safe, secure storage that Terremark has always been known for.

This document explains the architecture behind BLM, functionality and usage scenarios for customers to reduce their monthly Data Backup and Restore bill.

HOW BLM WORKS

Overview

BLM allows Terremark customers to reduce their monthly backup bill by moving data from the online storage to the BLM storage. The BLM storage is a near-line storage that slightly increases the time-to-restore but still allows data to be stored in Terremark's secure, fully-redundant infrastructure. BLM requests are generated (either on a scheduled basis or manually by the customer) to move data in the Terremark Data Center, from online storage to the BLM Storage on a periodic basis based on your corporate data retention policies. Once in BLM, you can access the data for searches and restores directly through the Web Portal or by calling the Terremark Help Desk.

If data needs to be restored from BLM storage, one or more restorable images are created that provide a complete snapshot of the data along with associated meta data. These restorable images are downloaded through the Web Portal or shipped on a portable disk. The last step of the process is to command the Terremark Backup Appliance (named Gateway) to access the restorable images on your local LAN or attached disk to perform a high-speed restore. Figure 1 provides a functional diagram.

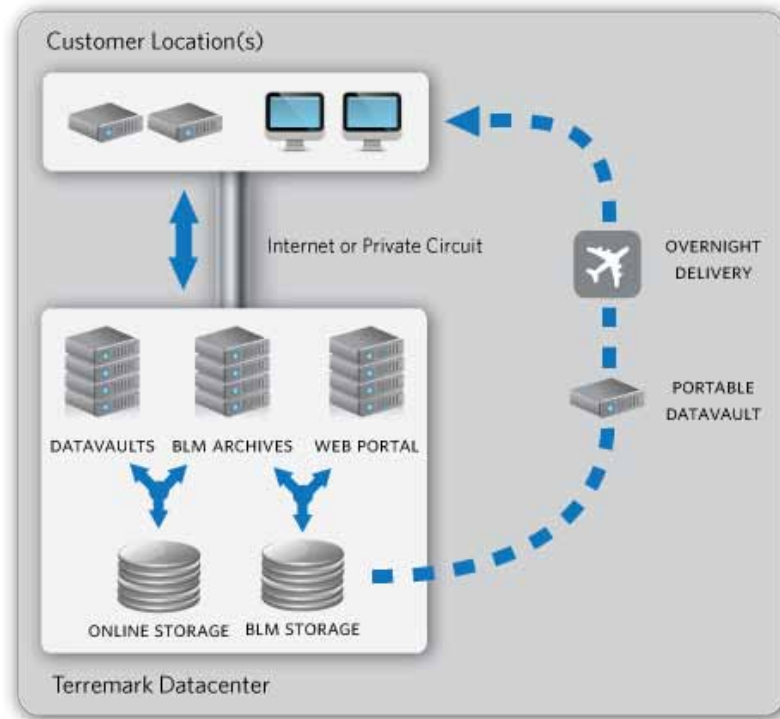


FIGURE 1 FUNCTIONAL DIAGRAM OF BLM

Data Buffers

Within the BLM Storage are two distinct data stores: the staging buffer and the consolidation buffer. As data transfers from the online storage, it is placed in the staging buffer (the short-term buffer) in an archive package. There can only be one archive package per backup set open at any one time. In the staging buffer archive package, the same number of files exists as in the online DS-System storage. After the archive package has been closed, the data is transferred to the consolidation buffer (the long term buffer). IN the consolidation buffer archive package, the source files (possibly in the millions) are compressed into very large (~100 GB) consolidation data files. The archive packages will remain in the consolidation buffer for a period of time that is determined by the company's data retention policies. Finally, as restorable images are created, data is copied to the restorable image buffer so that they may be downloaded or copied to a portable disk for overnight delivery to the customer. Figure 1 shows the architecture.

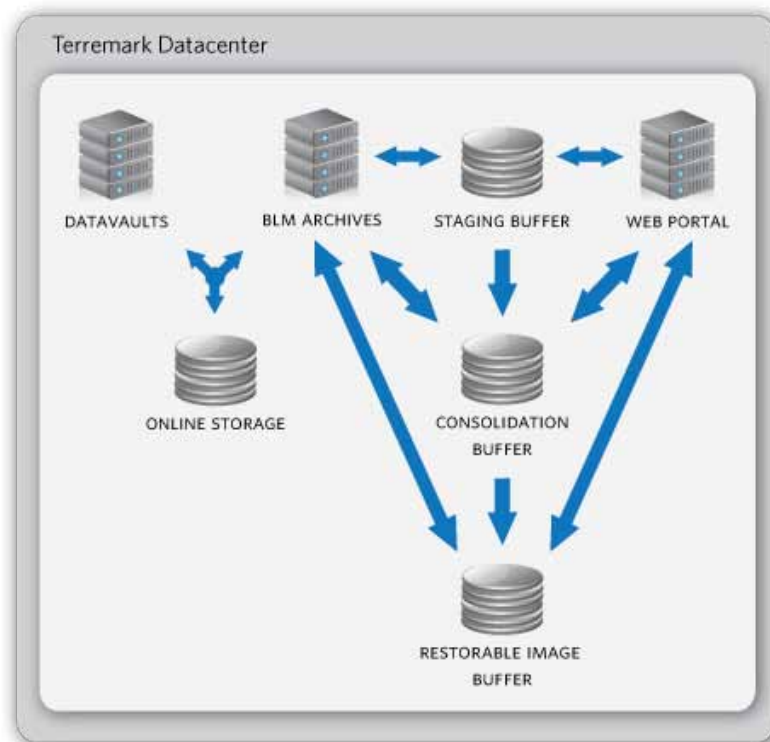


FIGURE 2 BLM STORAGE ARCHITECTURE

Archive Packages

Each archive package is for one backup set. A given backup set can have multiple archive packages but only one open archive package at a time. The open archive package resides in the staging buffer and can have multiple BLM actions write data to it. Only closed archive packages can be written to the consolidation buffer. An archive package can be closed when any of the following occurs:

1. It is manually closed
2. It is forced closed because another BLM action has been initiated and this BLM action was directed to use a new archive package
3. It is automatically closed because the size limit was reached

Web Portal

The Web Portal allows users to search and retrieve data that is either in the staging buffer or the consolidation buffer. A search is across the meta-data only (file name, email from/to addresses, etc) because the contents of each file or email is encrypted. Once the proper data is accessed, it can be added to a restorable image. The restorable image contains all data and meta-data needed to perform a restore. The restorable image can be downloaded via the Web Portal by the customer directly or copied to a portable device and shipped to the customer via standard overnight delivery.

Restore

Once the restorable image is in the customer infrastructure, the Terremark Gateway server is directed to read the contents of the restorable image and perform the restore in exactly the same fashion as if the data were in the online storage.

BLM Archiver GUI

The BLM Archivers in the Terremark Data Center can be accessed via a Java client that is operated by Terremark Help Desk personnel. All functions available through the Web Portal are available through the BLM Archiver GUI.

Data Destruction

At any step in the process, an archive package can be destroyed and the hard disk overwritten by random data. A data destruction request can be initiated by the customer directly (via the Web Portal) or by calling the Terremark Help Desk. There are at least three independent steps that must occur before data is destroyed in the BLM.

USAGE SCENARIO: STATIC DATA ARCHIVE

Many companies have volumes of static data that do not change over time. BLM allows for this static data to be moved to the cost-effective BLM archive. Since the data does not change over time, the one generation in the BLM archive is considered the archive copy.

The write-off to BLM is accomplished automatically from within the Gateway with a retention rule. A retention rule can define a number of days (say N) such that when source customer files have not been modified for N days, the files are moved to BLM from the online DS-System. This movement of data is accomplished during an "Enforce Retention" function which can be scheduled on the Gateway.

The static data archive usage scenario is by far the most popular method by which BLM is invoked.

USAGE SCENARIO: DELETE PUSH

With BLM, all deletions of data from the online storage have the option to copy or move the data to BLM rather than losing it. This is referred to as delete push. One common scenario where delete push would allow a company to save online storage costs would be the "Deleted Items" folder of Exchange mailboxes. During a MLR backup, all emails that are deleted by the user but not deleted from the Exchange server (due to either an Exchange policy or users not cleaning the "Deleted Items" folders) are backed up to the online storage. After some time, the emails are cleaned out of the Exchange "Deleted Items" folder, but the backed up data still exists in the online storage. By running a Delete Push to move to BLM all emails that are deleted on the Exchange server, all the "Deleted Items" emails are still backed up by Terremark, but are moved off of the online storage.