

VERITAS™



VERITAS NetBackup FlashBackup™

**For Lightning-Fast Backups of
Large File Servers**

V
E
R
I
T
A
S

W
H
I
T
E

P
A
P
E
R



Table of Contents

Introduction	1
Product Description	1
FlashBackup Features	2
FlashBackup Improves Backup Performance	3
FlashBackup's Restore Performance	3
Which Customers Benefit the Most?	4
System Requirements	4
Supported Operating Systems and File Systems	5
Ordering FlashBackup	5
FlashBackup Roadmap	5
FlashBackup Backup Operation Technical Overview	6
FlashBackup Restore Operation Technical Overview	7
Find More Information on FlashBackup	8

Introduction

While your critical enterprise server is up and running and performing its primary function, the VERITAS NetBackup FlashBackup™ feature can make a lightning-fast backup in a fraction of the time required by conventional backup methods, while still providing data integrity and individual-file restorability. This document provides an overview of FlashBackup, an optional feature of VERITAS' data center strength backup and restore application, NetBackup.

Product Description

FlashBackup is a separately priced optional feature of NetBackup that greatly reduces the amount of time needed to backup enterprise servers storing very large numbers of small files, such as NFS, e-mail, and file servers.

FlashBackup achieves remarkable increases in backup performance by massively reducing the number of disk I/Os needed during backups. FlashBackup does this by reading file information directly off the raw partition containing the files being backed up, bypassing the file system itself. This permits much faster backups for servers that store a large number of files. The more files being stored, the greater the potential backup performance improvement.

During a typical file backup, a server must perform a number of disk I/Os and system calls to locate and read each file (really the file header) to be backed up in the file system. Since directories, inodes, and data blocks are typically spread out across a disk, this results in a high level of scattered disk I/Os. FlashBackup reduces backup time by avoiding system calls and reducing disk I/Os. FlashBackup accomplishes this by creating a map of the file system's content and then writing the entire raw partition to tape as a single, very large file. While map creation imposes some additional overhead, writing the raw partition to tape maximizes tape throughput, resulting in a faster backup overall.

FlashBackup restores are also different than standard NetBackup restores. First, FlashBackup retrieves the map of the file system's content from the backup image. It searches the map to determine the location of the blocks of the files to be restored. FlashBackup can then extract the required blocks from the backup image and recreate the files to be restored.

FlashBackup Features

Several features distinguish FlashBackup from standard NetBackup and other raw-partition backup methods:

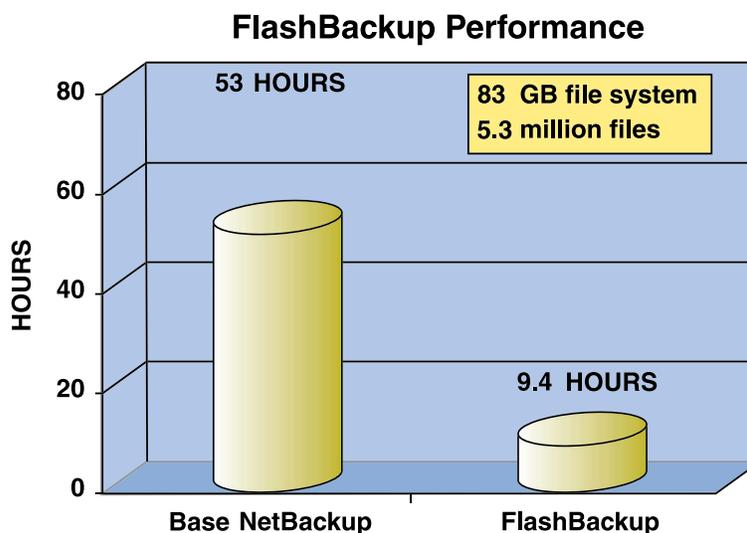
- Increased backup performance – When a server contains file systems that have a very large number of small files (thousands to millions) and most of the file system blocks are allocated, FlashBackup provides faster backup performance than other products on the market today.
- Individual file restore – FlashBackup provides a file system mapping capability that allows individual files to be restored from FlashBackup raw-partition backups. FlashBackup does not require a user to sacrifice granularity for performance.
- Raw-partition restores – FlashBackup can also restore entire raw partitions using the standard NetBackup mechanism.
- File systems remain mounted during raw-partition backups – It is not necessary to unmount the file system or mount it read-only, as with other raw partition backups. FlashBackup has an online backup capability that allows the file system to remain mounted on the raw partition, allowing it to be modified while the backup is in progress.

FlashBackup Improves Backup Performance

FlashBackup was developed to reduce the amount of time needed to perform a backup on servers that contain file systems with a very large number of small files. Performance gains observed in the field have been very impressive. In one documented installation, the backup time of a server with an 83 GB file system containing 5.3 million files was reduced from over 53 hours to a little less than ten hours, which amounted to a performance gain of over 500%.

During a backup, FlashBackup needs time to create a file system map and write the file system to tape. Assuming there are a large number of small files; FlashBackup will map approximately 4000 files a second. Using DLT technology, a file system can be written to tape at 8 MB a second. If we apply these guidelines to our example of the 83 GB file system with 5.3 million files, we expect a backup time of roughly 2.9 hours for mapping and 6.5 hours for writing to tape, which adds up to just 9.4 hours.

A number of variables affect FlashBackup performance, such as system load, machine resources, file system fragmentation, average file size, average file name length, and average number of files per directory. The numbers mentioned here are only meant to give you a rough idea of what FlashBackup's efficiency can be.



Source: Internet resume provider

FlashBackup's Restore Performance

During a standard restore, NetBackup retrieves the file to be restored from the backup image. With FlashBackup, the restore program must first retrieve information about the file from the file system map in the backup image, then retrieve the file's data blocks from the backup image. If the number of files to be restored is large and the file system was badly fragmented, it may be faster to restore the entire raw partition than a large number of individual files.

The time taken to retrieve the file system map is usually very short compared to the time taken to retrieve the required files. For this reason, the time needed for the FlashBackup restore will at most only be slightly longer than that of a standard NetBackup restore, when working with a small to moderate number of files.

As mentioned, FlashBackup is highly recommended for systems with a large number of small files. Since the data blocks that make up a small file are usually not scattered across the disk, performance-degrading fragmentation should not be a problem, as with other raw-partitioned backup applications on the market.

Which Customers Benefit the Most?

FlashBackup was specifically designed to increase backup performance for servers with file systems containing a large number of small files. FlashBackup can greatly reduce the time required for a complete backup when deployed on a server with a file system containing thousands to millions of files and 70% or more of its storage capacity used up. However, a server that is only 20% full, with file systems containing a small number of large files, will not see a large performance increase, but may benefit from FlashBackup's ability to perform consistent point-in-time backups of live file systems.

Some highly successful installations of FlashBackup with VERITAS NetBackup include ISPs with large e-mail servers, and a group of petroleum and geological research organizations with large file servers. Any organization with servers that harbor a very large number of small files are a very good fit for this solution.

System Requirements

Mapping a large file system demands sufficient CPU cycles, RAM, and temporary disk space. However, the system requirements cited below are much lower than what is needed for hardware split-mirror solutions that can provide the basis for backup operations. Hardware split-mirror solutions require 3 times the disk space and can only restore entire volumes.

- RAM: approximately 24 times the number of millions of files in the file system. For example, approximately 48 megabytes of memory are needed for a file system with 2 million files.*
- Disk space: approximately 64 times the number of millions of files in the file system. Therefore, approximately 128 megabytes of memory are needed for a file system with 2 million files.*
- CPU: FlashBackup will use about half of a CPU's cycles for each concurrent backup mapping operation. This should be borne in mind if FlashBackup will run on a system that performs other operations. If you plan on running multiple concurrent instances of FlashBackup, you need to add up the requirements for all file systems concerned.

*RAM and Disk space requirements will be lower for environments with thousands of files

Supported Operating Systems and File Systems

FlashBackup is currently supported on Solaris and HP-UX servers running UFS, JFS, or VxFS file systems. Please contact your VERITAS sales representative to obtain an updated list of currently supported platforms.

Ordering FlashBackup

FlashBackup is a separately priced, optional feature of NetBackup. Each server using FlashBackup functionality requires a license. In addition, at least one media kit must be purchased. Please consult a VERITAS sales representative or reseller for pricing information.

FlashBackup Roadmap

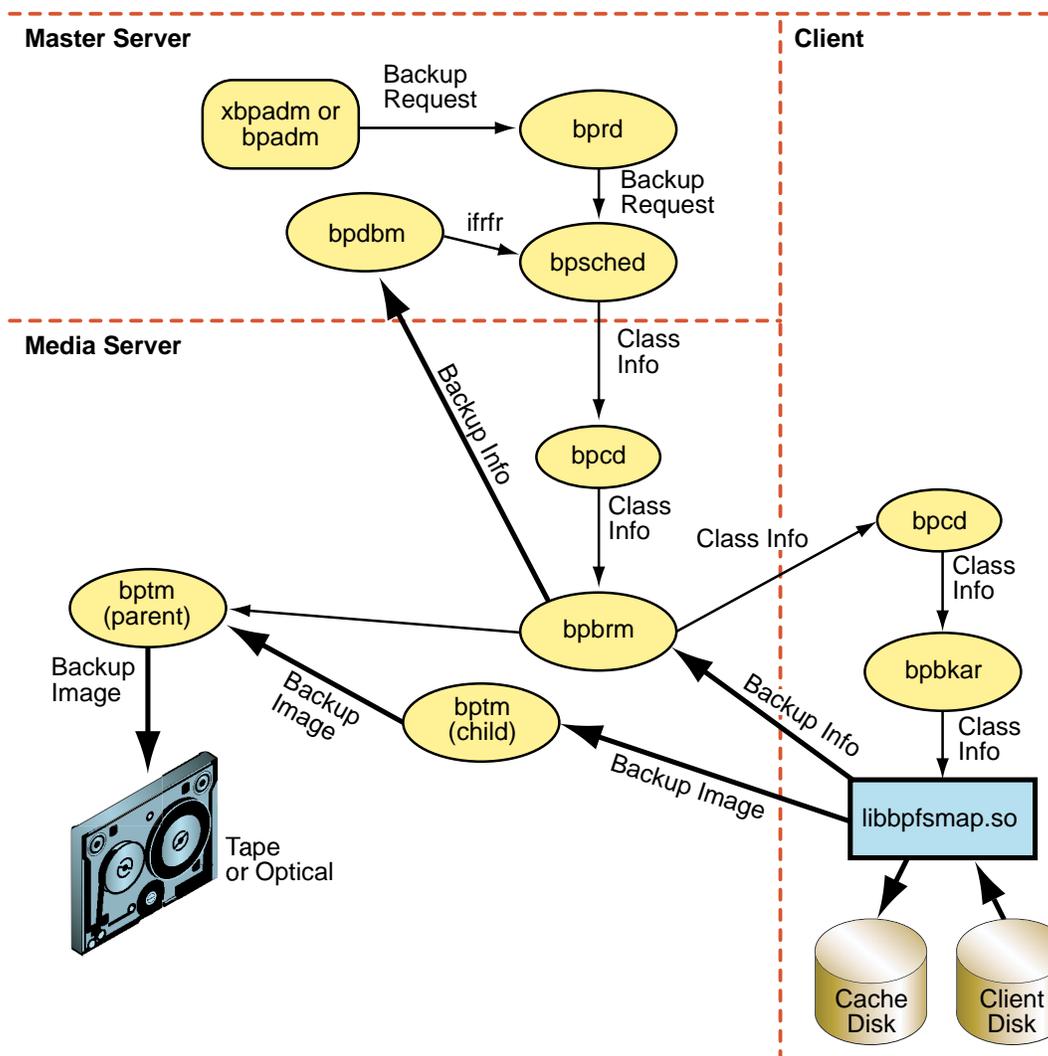
FlashBackup technology is currently being developed for both the Windows 2000 and AIX operating systems.

FlashBackup Backup Operation Technical Overview

Scheduled backups and immediate backups using FlashBackup are initiated and processed in the same manner as standard NetBackup. However, the individual file restore from raw attribute that defines FlashBackup, causes bpsched to send a special flag to bpbm along with the other class information. Using the special flag, bpbm starts bpbkar, which causes bpbkar to invoke the FlashBackup dynamic shared library, libbpfsmmap.so. During the backup, libbpfsmmap.so performs the following tasks:

- Creates a file system map that represents files contained in the raw partition
- Creates a bit map that associates blocks belonging to selected files with their physical addresses in the partition
- Turns on the raw-partition snapshot mechanism so that when a block is updated in the partition, the previous contents of the block are written to the cache partition
- Sends the backup information to bpbm, which in turn directs the backup information to bpbkar
- Transmits the backup image to bpbm
- Turns off the snapshot mechanism

The following figure shows the major NetBackup components involved on a FlashBackup backup operation.



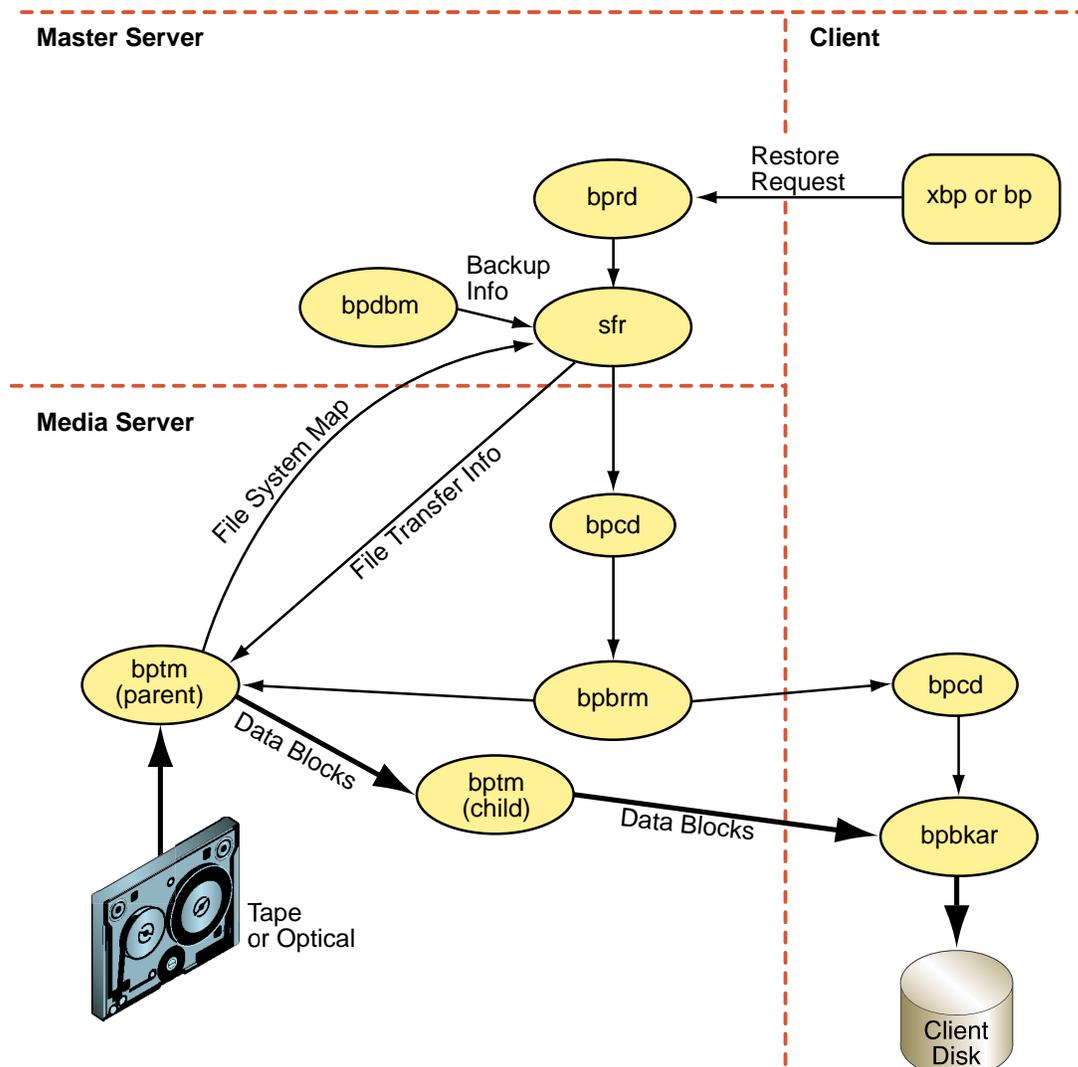
FlashBackup Restore Operation Technical Overview

During a restore, the user browses for and selects files from FlashBackup images in the same manner as standard backup images. The difference is that bprd calls the sfr program when processing FlashBackup images. During a restore, the sfr program retrieves information about the backup from bpdbm. The backup information is composed of:

- File system map name and its media location
- Bit map name, and its media location
- Raw-partition name, and its media location

Using the backup information, sfr retrieves the file system map by directing bptm to the exact location on the tape. A similar procedure is followed to obtain the bit map data when restoring from incremental backup images and to obtain the raw-partition data when restoring the raw partition. The sfr process then uses the information contained in the file system map to direct bptm to the location of the individual file data blocks on the tape. bptm then reads the file data blocks from the tape and writes them, via tar, to the destination disk location.

The following figure shows the major NetBackup components involved in restoring files using FlashBackup.



Find More Information on FlashBackup

- **FlashBackup System Administrator's Guide** – Explains how to install, configure, and manage FlashBackup.
- **NetBackup Release Notes - UNIX** – Explains how to install NetBackup for UNIX server and client software and provides other important information, such as supported platforms and operating systems
- **NetBackup System Administrator's Guide - UNIX** – Explains how to configure and manage NetBackup on a UNIX platform
- **NetBackup User's Guide - UNIX** – Explains how to perform user-directed backups, restores, and archives on a UNIX client
- **Media Manager System Administrator's Guide** – Explains how to configure and manage the storage devices and media that UNIX NetBackup servers use for backups
- **Media Manager Device Configuration Guide** – Provides information about configuring storage devices on UNIX systems
- **NetBackup Troubleshooting Guide** – Provides troubleshooting information for all NetBackup products.



V
E
R
I
T
A
S

W
H
I
T
E

P
A
P
E
R

VERITAS Software Corporation
Corporate Headquarters
1600 Plymouth Street
Mountain View, CA 94043
650-527-8000 or 800-327-2232

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com
e-mail us at sales@veritas.com