

De-Duplication of Files for Remote Storage

© Copyright 1998-2008 American Megatrends, Inc.

All rights reserved.

American Megatrends, Inc.

5555 Oakbrook Parkway, Building 200

Norcross, GA 30093

TRADEMARK AND COPYRIGHT ACKNOWLEDGMENTS

This publication contains proprietary information that is protected by copyright. No part of this publication can be reproduced, transcribed, stored in a retrieval system, translated into any language or computer language, or transmitted in any form whatsoever without the prior written consent of the publisher, American Megatrends, Inc. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. American Megatrends, Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

FOR ADDITIONAL INFORMATION

Call American Megatrends at 1-800-246-8600 for additional information. You can also visit us online at ami.com.

LIMITATIONS OF LIABILITY

In no event shall American Megatrends be held liable for any loss, expenses, or damages of any kind whatsoever, whether direct, indirect, incidental, or consequential, arising from the design or use of this product or the support materials provided with the product.

LIMITED WARRANTY

No warranties are made, either express or implied, with regard to the contents of this work, its merchantability, or fitness for a particular use. American Megatrends assumes no responsibility for errors and omissions or for the uses made of the material contained herein or reader decisions based on such use.

DISCLAIMER:

Although efforts have been made to assure the accuracy of the information contained here, American Megatrends expressly disclaims liability for any error in this information, and for damages, whether direct, indirect, special, exemplary, consequential or otherwise, that may result from such error, including but not limited to the loss of profits resulting from the use or misuse of the information contained herein (even if American Megatrends has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to American Megatrends at the address shown on the back cover of this document.

American Megatrends provides this publication "as is" without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a specific purpose. Some states do not allow disclaimer of express or implied warranties or the limitation or exclusion of liability for indirect, special, exemplary, incidental or consequential damages in certain transactions; therefore, this statement may not apply to you. Also, you may have other rights that vary from jurisdiction to jurisdiction. This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. American Megatrends may make improvements and/or revisions in the product(s) and/or the program(s) described in this publication at any time.

Introduction

Data deduplication (often called "intelligent compression" or "single-instance storage") is a method of reducing storage needs by eliminating redundant data. Data deduplication offers other benefits. Lower storage space requirements will save money on disk expenditures. The more efficient use of disk space also allows for longer disk retention periods, which provides better recovery time objectives (RTO) for a longer time. Data deduplication also reduces the data that must be sent across a WAN (Cloud Storage) for remote backups, replication, and disaster recovery. Data deduplication is often used in conjunction with other forms of data reduction such as conventional compression and delta differencing.

Abstract

A fast efficient, less CPU intensive de-duplication method for files is presented. The invention is based on the simple fact that the probability of nth version of a modified file will have more duplicates with the (n-1) Th version, as compared to other previous versions. The algorithm presented is a single pass algorithm, which also ensures that the time taken to de-dupe is always on the order of n, where n is the file size. The resulting de-duped files can be better used in situations where the backup storage is in remote locations, as in online cloud storage. Also, the dependency on previously backed up data is very less, hence reducing the cost of accessing backup storage as in the case of online cloud backups.

Need for Data Deduplication

Data protection techniques such as CDP are becoming more common with the reduction in the cost of storage. De-duplication and compression are techniques that are used in such CDP solution for optimizing the backup storage used. With the arrival of online cloud storage, increasingly large numbers of CDP data are being backed up into such remote storage. The reason being obvious that the cloud storage claims very high reliability on the backed up data. But such online cloud storage also throws in challenge to the CDP software to perform de-duplication. A CDP process would normally require having access to the backup data and its checksum to remove duplicates. This would mean that the CDP software running in the data source should have access to the data in the cloud storage. That would mean a large amount of data transfer from the cloud storage to the data source system. This would normally restrict the source level de-duplication when using cloud storage. But cloud storage would be the preferred method of storage for mobile computers where the de-duplication is best done at the source – the mobile computer itself.

Efficient Data De-Duplication

AMI's advanced de-dupe logic claimed here uses a single pass algorithm along with an efficient matching logic which makes the de-duplication to happen quicker. Along with this, de-duplication is always done with previous version of file (which will be available locally) to reduce the dependency on large master file which may sometimes be stored in the remote system.

Each and every de-duplication is done with the previous version and hence the new version will have pointers to previous version and also to new blocks which are newly added in this version. The result will be creation of the following:

Master file – those blocks that are new and are of de-dupe size (in our case, this is 512 bytes)

Data file – those blocks that are new and smaller than de-dupe size (in our case, < 512 bytes)


Version Map – a final de-dupe map which can be used to regenerate the original file – this will have pointers to Master file, Data file and pointers to previous version file.

Since de-dupe is done with previous version file, the previous version pointers are actually made against the Master file of previous version. A Version Map of version X may have pointers to Master file (X) and also to Master file (X-1). This may also chain back to as many versions as possible – if that particular block never got modified after a very old version.

While this dependency helps in providing an efficient de-duplication, it also brings in a difficulty that a 100 version file might depend on the Master file of version 1. Further, there may be only one block in the version 1 Master file that is being referenced and the rest are obsolete. This will result in poor space reclaim during version retention. Though the version retention is set to say 10, we might end up keeping all the 100 versions of Master file. To avoid this, we bring in the concept of **reference resetting**.


By *reference resetting*, after every n versions, say 4, when a new version of file is available, instead of de-duping with previous version, we will consider it as a new file. By this, the 5th version file becomes new in a chain, while versions 1, 2, 3 & 4 may have reference to master files 1, 2, 3 & 4. Version 5 will never refer to any of the previous versions and hence the versions 6, 7, 8 and so on will not refer to versions 1-4. By resetting the reference, we make them as individual chain and hence when the retention deletes version 4, we are safe to delete all the previous version Master files. But until version 4 is not deleted, we may not be able to delete master files of versions 1-3, since version 4 may have reference to them.

This resetting should not necessarily be at the boundary of the version limit and hence makes it easy for the user to change the version limit as many times as he/she prefers.



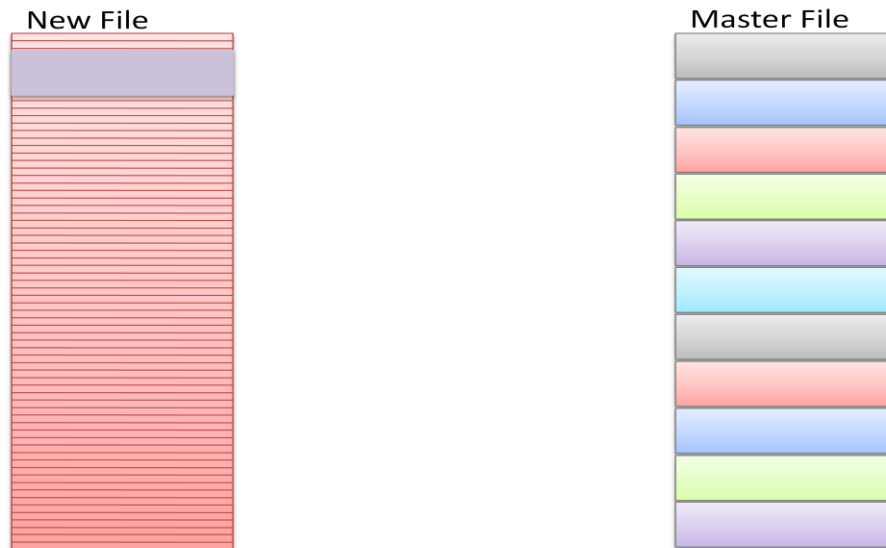
DATA DE-DUPLICATION

The above figure represents the new file – second version and the first version master file. Each tiny rectangle in the new file represents the slide size. Since it is one byte in our case, the complete file will be a multiple of the slide size. The right side represents the first version master file, each colored rectangle representing a fixed sized block, here 512 bytes. The master file will always be a multiple of block size. The first block add-with-carry in the new file does not match with any of the blocks in the existing master file, which is denoted by the non-matching colors (very light red as against the vibrant colors on the right side).




DATA DE-DUPLICATION

The above figure represents the algorithm slides on the new file by slide size, the add-with-carry of this block also does not match with of that in the right side. The new add-with-carry was computed by subtracting the first byte and adding the new byte to the existing add-with-carry. This addition and subtraction are less CPU intensive, single-pass process.



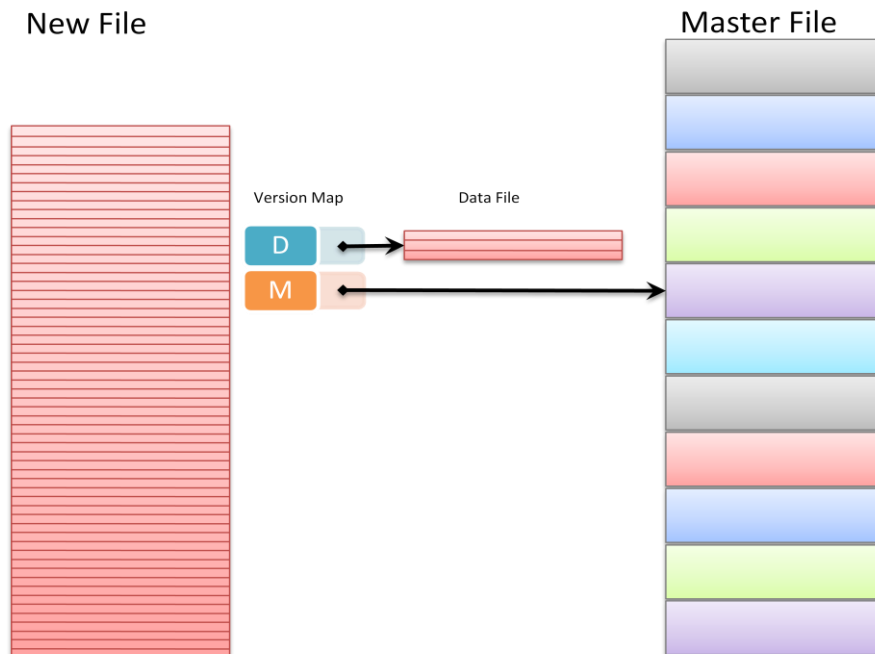
DATA DE-DUPLICATION

This represents the algorithm slides on the new file by slide size; the add-with-carry of this block also does not match with of that in the right side.



DATA DE-DUPLICATION

The above figure represents yet another slide on the new file by slide size, but this time we see that the add-with-carry on the new file matches with the add-with-carry of the red block on the right side. We will read the data from the master file and compare it to ensure that the data matches.



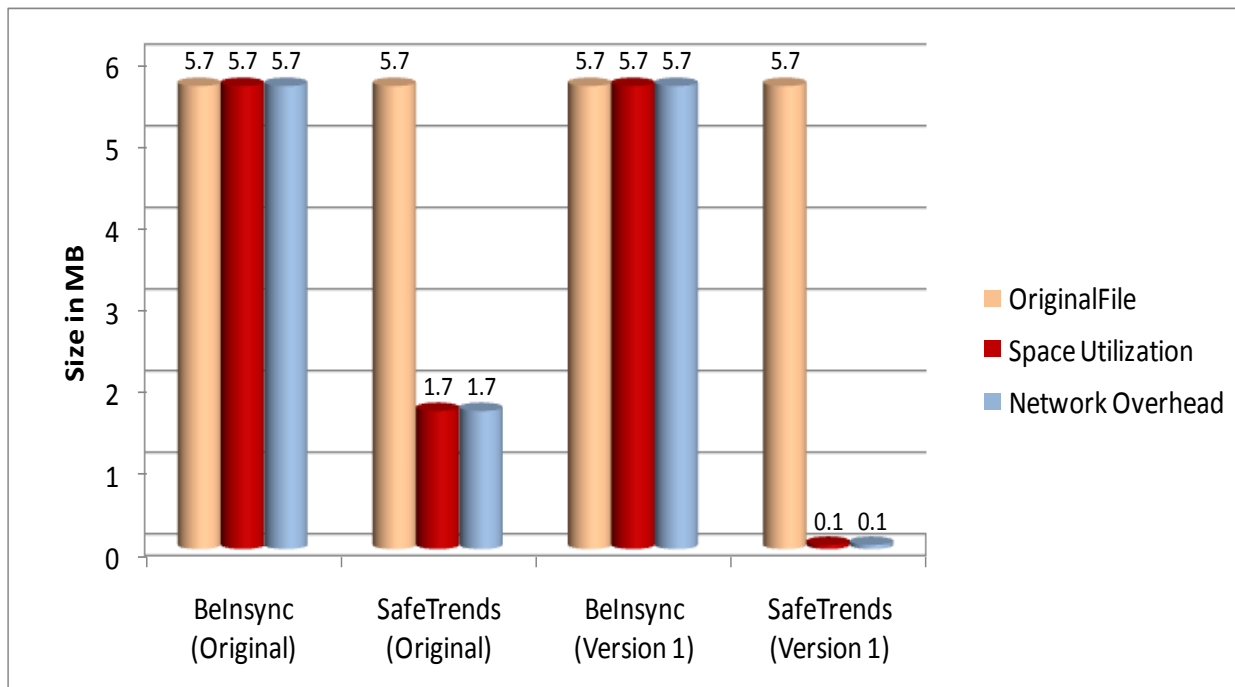
DATA DE-DUPLICATION

In the above figure we see that the data also matches. Now the existing non-matching data above the matching segment is added to the data file. A version map “data file entry” to that effect is added. Also a “master file entry” is added to point to the existing master file. This block is now de-duped and the duplicate removed.

By proceeding as above, the entire file is de-duped creating new version map, data file and master file entries.

How does it benefit your business?

- Data deduplication reduces the amount of disk space you need for backup by 90% or more. That means dramatic savings in space, power, cooling and overall costs. With deduplication, you keep much more backup data on fast-recovery RAID arrays than with any conventional disk backup system. You get fast, reliable backup and restores, increase the number of recovery points, and yet use less disk capacity.
- Data deduplication reduces the network bandwidth needed for replication by 90% or more. Deduplication provides automated DR and site-loss protection using existing networks, reduced media handling (especially in distributed environments), and lower costs. It finally makes WAN-based replication a practical solution for DR protection.



The above chart was a comparative study and benefits of backing up two versions of a 5.7 MB Word Document.

This Data Deduplication, being a single-pass, light weight less CPU intensive yet producing good amount of de-duplication, this would be the preferred method for CDP process to store data in online cloud storage or remote storage. The fact that it is less CPU intensive comes in handy for usage in mobile computers that require data protection.

Why AMI?

AMI offers a wide array of disaster recovery and high availability solutions for your business needs. We provide services that range from storage needs analysis to the design and implementation of a custom disaster recovery solution. We can help your business plan for when things are at their worst, while reducing the cost and complexity of your storage environment. For more information on AMI StorTrends solutions, visit www.StorTrends.com, email to sales@ami.com, or call (800) U.Buy.AMI.