BACKUP STORAGE BLOCK-LEVEL DEDUPLICATION WITH DDUMBFS AND BACULA

Heitor Faria[1], Jacir Luiz Bordim[2] and Priscila Solis Barreto[3]

[1]MSc in Applied Computing (MScAC), Brazil
[123]Department of Computer Science (CIC), University of Brasilia (UnB), Brazil

ABSTRACT

Data deduplication is not only becoming an essential backup system component, because it reduces storage space requirements, but also a critical one, since the performance of all the backup operation depends on storage throughput. This paper is focused on the backup data volume ratio of deduplication when stored in file systems with this feature, and how a newly developed volume format can improve it. Bacula backup software is currently the third most popular backup software worldwide and it's also open source, therefore selected to this performance analysis study. Corresponding to the results, it is possible to verify that the proposed Bacula aligned backup format brings gains of 82%, even for a workload that emulates a relatively small backup policy and with minimum impact on backup or restore times.

KEYWORDS

Backup volume, deduplication format, file system, data reduction

1. INTRODUCTION

According to Gartner[1] Consultancy Report by Pushan Rinnen and Russel (2015), digital data has snowballed to a level that frequently leads to backup storage capacity depletion, and it's imperative to deploy mechanisms to reduce backup software storage demand, without sacrificing data retention or downgrading policies.

In this way, deduplication is one of the most researched fields in information technology. Nowadays, there are several [3] software and hardware options to deploy storage block-level data redundancy elimination techniques.

Deduplication is expected to reduce significantly backup data occupancy [4], since jobs just repeat the copy of basically the same files over a fixed schedule. Duplicated files backed up from the same host (e.g.: an email sent to hundreds of corporate users), would also be stored only once in the backup storage. The amount of reduced redundant information may vary according to enterprise workload, that can be very large.

However, unprepared property and even open source software write backup jobs in its own format, causing poor deduplication performance. As stated by Sibbald [5], most backup software archive programs stream the metadata, headers, and file data to a single file compacting and mixing them together.

As shown in Figure 1, different servers backup jobs store data in small chunks interspersed with header and metadata, which makes the volumes unique and difficult to deduplicate.

DOI : 10.5121/ijait.2017.7401

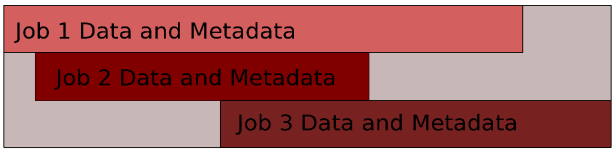

Figure 1: Backup software traditional volume writing format.

The purpose of the present work is to verify Bacula[6] backup software currently community version deduplication gains when storing backup volumes in a speed focusing deduplication Linux filesystem, ddumbfs[7], and comparing to a new developed deduplicative Aligned Volume format to be patched to its code still this year.

For result validation purposes, Bacula old and new aligned performance results are compared to another third party developed open source deduplication backup software named ZBackup[8].
In Section 2 we present the State-of-the-Art. In the Section 3, the Related Work. Section 4 shows the Methodology. In Section 5, the Results. Finally, Section 6 draws some conclusions and final remarks. Section 7, Future Works.

2. STATE-OF-THE-ART

As reported by Xia et al. [4], data deduplication is a dictionary based data reduction approach, due to its ability to effectively reduce backup storage or archiving datasets size by a factor of 4-40X.

In consonance with several authors [9, 10, 11, 12], the chunk-level data deduplication technique usually splits data blocks of a data stream into multiple data chunks of average size 4K or 8K. Each block is uniquely identified and duplicate-detected by a SHA-1 or MD5 hash signature, also called a fingerprint.

Pursuant to Constantinescu et al. [10], the secure fingerprint-based deduplication eliminates redundancy at the chunk or file level and, therefore, scales better than the traditional LZ77 and Huffman GZ coding based compression.

To store backups, Bacula is the open source software used in this paper. Conforming to Google Trends[13], it's the third most popular enterprise server backup system worldwide. Full, differential, incremental file level multiplexed jobs [14] are supported by Bacula. Backups are written into tapes or disk files, using open format bytes-sequence volumes.

The new proposed Bacula deduplicatable archive volume is described by Sibbald [5] as a data container from an original stream of multiple files, and which can be optimally deduplicated by an underlying deduplication storage system.

As displayed in Figure 2, the method comprises receiving data records representing metadata and file data, at least a part of which are already separated, separating the metadata and the file data into a first file and a second file, the first file and the second file being paired, the first file called Metadata Volume containing metadata, header data and references to the file data, and the second file called Aligned Volume containing file data only.

Figure 2: Bacula new aligned volume writing format.

A further part of the records which contain both metadata and file data are separated into metadata and file data and then subjected to the step of separating the metadata and the file data into the first file and the second file.

3. RELATED WORK

In line with Grant [12], data-protection class workloads including backup have seen a strong shift from tape-based to disk-based platforms, but little analysis of the characteristics of backup workloads and its challenges has been published.

Also, while tapes have sequential access, disks are random accessed and requires significantly less effort to have stored data deduplicated. Still, the authors stated that backup storage workloads have inherent high degree of data redundancy [12, 15], and high throughput make deduplicating techniques important.

Consistent with Zhang et al. [16], Bacula is an open source backup software whose copied servers data is stored in the storage server with no index. Even for disks, read and write operations are made in similarity with a tape library, when each recovery or backup is sequential. There is no distinction of disk or tape volumes from the storage device point of view and multiple server data of different platforms can be stored in the same storage volume. The backed up files metadata is also stored in a SQL database for quick restore job files selection.

Currently there are several deduplication file systems nowadays, such as lessfs[17], opendedup[18] and others [19, 20], however ddumbfs was chosen for being both open source and focused on faster operations thanks to its very simple index design, which is very important for shorter backup windows.

The ddumbfs aims to be a fast in line deduplication filesystem for Linux, based on FUSE and released under the GNU General Public License. Before to writing a block of data, ddumbfs check if any identical block was already saved on disk. It generates the SHA1 or the TIGER hash of the block and search the index for a similar one. If there is a match, the reference to the already stored block is used. Else, the block is saved as a new one and the index is updated.

ZBackup deduplicates data from different tar[21] streams and between them, using a different deduplication technique. It uses a sliding window with a rolling hash over the input at a byte granularity and checking whether the block in focus was ever met already. If a rolling hash matches, an additional full cryptographic hash is calculated to ensure the block is indeed the same. ZBackup does not need an underlying deduplication system storage.

4. METHODOLOGY

Most performance problems are unique and deserve a specific and systematic approach [22]. The metrics, workload, and evaluation techniques used for one problem generally cannot be used for the next problem. However, there are steps common to all performance evaluation projects that help to avoid analysis mistakes.

There are 10 steps proposed by Jain[22] and presented in Figure 3. The first one consists of stating goals and boundaries of the evaluated systems, and the others should lead to presented adequate results. However, it is stated the cycle can be restarted if the results are not consistent, suggesting it is a not fixed flow. This will be the methodology used in this project.

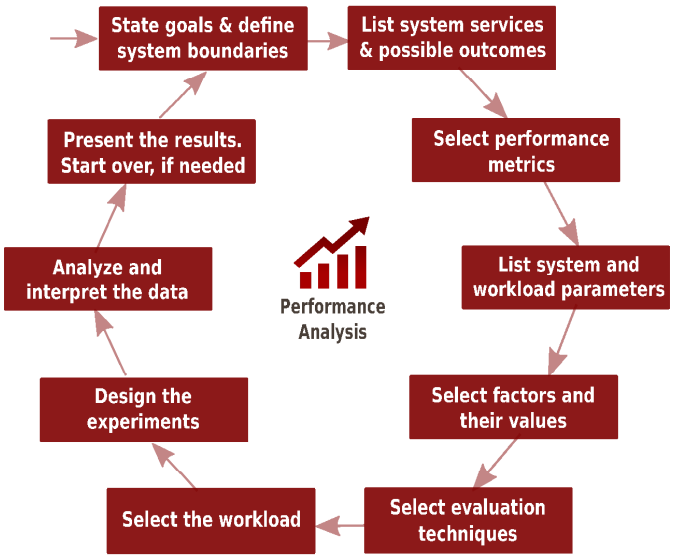

Figure 3. Steps for a correct performance analysis.

The goal of the experiment is to determine stored backup volumes deduplication performance between Bacula Community (version 7.4.4) and Enterprise[23] (version 8.8.2), since the last one contains the implemented new Aligned Format code. Version 1.1 of ddumbfs was used as well as the ZBackup source code version downloaded from Git [8] in May 12th 2017.

In both Bacula Editions their three services, Director, File (Client), Storage daemons, and also the Catalog database, runs in a local physical machine. ZBackup is also tested in the same environment.

The performance metrics are:

- the backup size in comparison with the stored deduplicated data;
- the deduplication efficiency;
- backup jobs duration;
- backup execution throughput;
- the restore time of a full backup.

Experiment is conducted using a i5-6400 Intel CPU at 2.70GHz, 16GB RAM desktop. A Kingston SV300 SSD hosts the Ubuntu 17.04 operating system and applications. Another Samsung EVO 850 500GB SSD stores backup data in a ddumbfs mounting point, and also ddumbfs index, which is very recommended to avoid deduplication filesystem transactions bottleneck. ZBackup tests uses the same infrastructure for its archiving area.

Real workload data is used, where a full backup job contains mixed type of data: Linux Ubuntu operating system files (22.74GB), a folder containing running virtual machine disk images (9.8GB) and a text database dump (8.3GB), totalizing 640,844 files and 40.84GB bytes of information to be copied.

Software compression is not used for this comparison since it does not deduplicate well due to the tendency of randomizing data. Series of five full backup jobs and four differentials is submitted, simulating the retention of one month using grandfather-father-son (GFS) rotating scheme (five weekly full and four daily differential backups).

All the results have 95% confidence interval, with three times of presented experiment unit results executed to achieve it. 15 full, 12 differential backup and 15 restore jobs were run for Bacula Enterprise plus ddumbfs, Bacula Community plus ddumbfs, and for Zbackup.

5. RESULTS

As shown in Figure 4, there is no benefit when storing just one full backup job using the new aligned volumes (Job1, with 41GB size). However, deduplication gains scales greatly when more jobs are executed, since most data does not need to be copied again. Deduplication storage size remains almost stationary, and even for smaller differential backups (Jobs 6-9) the economy is significant. The new solution should bring economy gains for most of enterprises disk based backup systems.

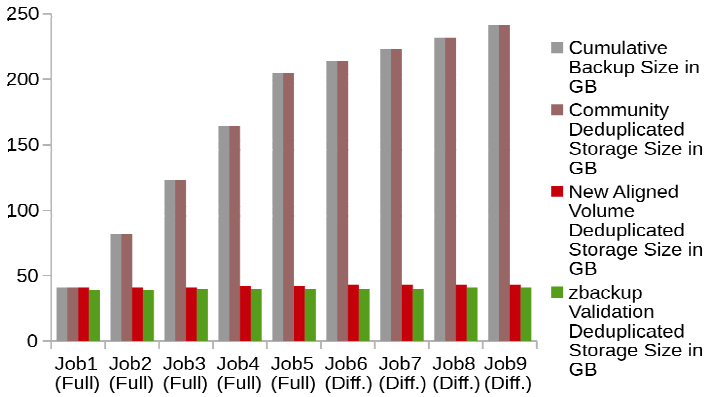

Figure 4: Backup size in comparison with the stored deduplicated data.

ZBackup first backup job had a few gains due to different sliding window technique, with 39GB of size, just approximately 5% lower. This difference might be acceptable in face of another result metrics trade-off.

As exhibited in Figure 5, old Bacula volume format is basically undeduplicable, while new technique has 82,16% of storage space saving for the current workload. This value would be higher if more backup jobs were executed. The result is similar to ZBackup measured efficiency of 83%.

It is important to remember that software compression was not enabled for the Bacula Community version and it would alternatively bring a 56% efficiency for the same workload using GZIP6 algorithm, approximately. However, compression would present slower backups, the economy gains doesn't scale like in deduplication, and this comparison isn't in the scope of this paper. Compression is also disabled in ZBackup configuration for equality of conditions.

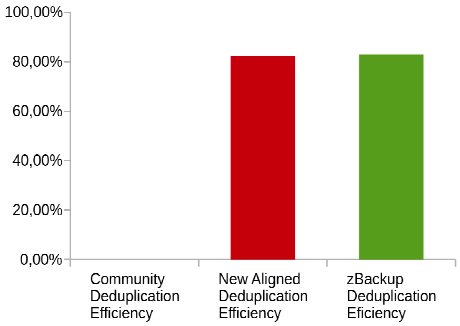

Figure 5: Deduplication efficiency.

As advertised in Figure 6, deduplication matches have definitively an impact in the backup duration, although it is relatively small. While average undeduplicable Backup Jobs finished in 295 seconds, new aligned backups terminated 322 seconds. The observed difference is around 9%.

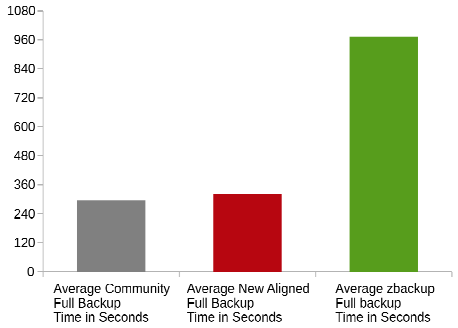

Figure 6: Backup jobs duration.

ZBackup average backup job duration was 973 seconds, almost 302% slower than the new Bacula Aligned format backup. In fact, the first full backup of the deduplication engine took 3186 seconds, but the next ones are much faster.

In the same way, as exposed by Figure 7, the deduplicative volume format also suffered a minor impact when demanded, while Bacula Community presented 140.78MB/s of throughput, against 125.99MB/s of the new technique. In relation to ZBackup, Bacula Aligned Volumes was nearly 314% faster.

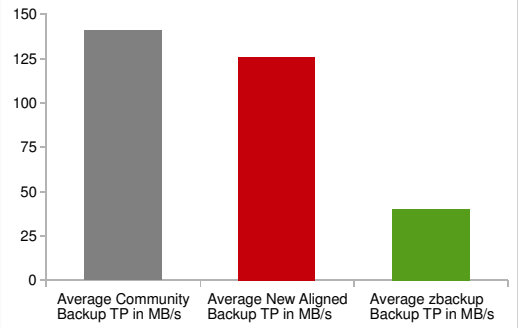

Figure 7: Backup execution throughput.

At least, Figure 8 presents a slightly poorer deduplicated volumes performance when presenting data for restore. A full Restore Job was executed under a 211 seconds time frame using the undeduplicable volumes, but 258 seconds were spending restoring from the deduplicated volumes (about 22% slower).

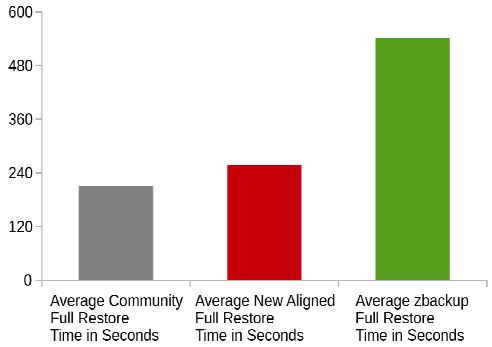

Figure 8: The restore time of a full backup.

Backup and restore jobs were slower when storing backups in the deduplicative filesystem, what would suggest a trade-off for the technique. However, an indexing bottleneck could be caused by using an equally fast SSD for deduplication engine index and data storing. Repeating this experiment storing backups in traditional slower HDD might present results where there is no difference between deduplicative and non-deduplicative Bacula volumes job durations.

ZBackup restore was 210% slower, showing both Bacula Community and Enterprise are more suitable for larger backup workloads.

6. CONCLUSION

Deduplication techniques are widely studied, but developers must be aware data must be stored in a compatible format, otherwise they might be considered useless.

Also, backup workload tends to be excessively large but very repetitive, therefore an ideal object for deduplication. Not only stationary data but also running virtual machines disks or databases can fetch gains while backed up to a deduplicative device.

For this paper experiment, the deduplication efficiency was of approximately 82%, with tests that emulates one month of full weekly and daily differential backups. It is possible to see clear trends for the presented results, meaning the workload is significant and that stochastic process can be very accurate forecasting the size of future deduplicated backup jobs. In this way, policies with more frequent or bigger retention times can drive to an even higher measured efficiency.

There is a minor impact in backup and restore duration, but it is an acceptable trade-off. The new Aligned Format proves to be a good storage cost reducing new Bacula Community feature, and to be much more efficient than ZBackup in terms of backup and restore speeds.

More than ever, disk backups are becoming a feasible replacement for tape libraries, since deduplication is not a feature that can currently be efficiently deployed on the sequential magnetic media.

7. FUTURE WORKS

Bacula Aligned Format code already runs in production environment of large Data Centers worldwide [24], deployed in the Bacula Enterprise Edition software. The next step is porting this code to the Community version, making it available to the general public usage, study and distribution.

ACKNOWLEDGMENT

Heitor Faria would like to thank his UNB Master Degree classmates, staff, teachers, coordinator (Dr. Marcelo Ladeira), Wanderlei Huttel for helping him finding ddumbfs, Kern Sibbald for the paper idea and for always being supportive.

REFERENCES

[1] "Technology Research | Gartner Inc." [Online]. Available: http://www.gartner.com/ technology/home.jsp

[2] Pushan Rinnen and D. Russel, "Challenging Common Practices for Backup Retention," Gartner, Inc., USA, Tech. Rep. G00278794, Jul. 2015.

[3] M. Fu, D. Feng, Y. Hua, X. He, Z. Chen, W. Xia, Y. Zhang, and Y. Tan, "Design Tradeoffs for Data Deduplication Performance in Backup Workloads." in FAST, 2015, pp. 331–344. [Online]. Available: https://www.usenix.org/sites/default/files/conference /protected-files/fast15 slides fu.pdf

[4] W. Xia, H. Jiang, D. Feng, and L. Tian, "Combining Deduplication and Delta Compression to Achieve Low-Overhead Data Reduction on Backup Datasets." IEEE, Mar. 2014, pp. 203–212. [Online]. Available: http://ieeexplore.ieee.org/document/ 6824428/

[5] K. Sibbald, "Creating a universally deduplicatable archive volume," Patent, 2016, classificao internacional G06F17/30; Classificao cooperativa G06F17/30156, G06F17/30073, G06F17/30194. [Online]. Available: http://www.google.com/patents/ US2016005516

[6] "Bacula | Open Source Backup, Enterprise ready, Network Backup Tool for Linux, Unix, Mac, and Windows." [Online]. Available: http://blog.bacula.org/

[7] "ddumbfss homepage ddumbfs 1.1 documentation." [Online]. Available: http:// www.magiksys.net/ddumbfs/

[8] "ZBackup." [Online]. Available: http://zbackup.org/

[9] A. Muthitacharoen, B. Chen, and D. Mazieres, "A low-bandwidth network file system," in ACM SIGOPS Operating Systems Review, vol. 35. ACM, 2001, pp. 174–187. [Online]. Available: http://dl.acm.org/citation.cfm?id=502052

[10] C. Constantinescu, J. Glider, and D. Chambliss, "Mixing Deduplication and Compression on Active Data Sets." IEEE, Mar. 2011, pp. 393–402. [Online]. Available: http://ieeexplore.ieee.org/document/5749497/

[11] B. Zhu, K. Li, and R. H. Patterson, "Avoiding the Disk Bottleneck in the Data Domain Deduplication File System." in Fast, vol. 8, 2008, pp. 1–14. [Online]. Available: https:// www.usenix.org/legacy/events/fast08/tech/full papers/zhu/zhu html/

[12] W. Grant, "Characteristics of Backup Workloads in Production Systems," 2012.

[13] "Google Trends: Arcserve, Bacula, Netbackup popularity search worldwide," 2016, 1. [Online]. Available: https://www.google.com.br/trends/explore?date=all&q=arcserve, bacula,netbackup

[14] K. Sibbald, "Main Reference," 2011. [Online]. Available: http://www.bacula.org/7.4.x-manuals/en/main/Main Reference.html

[15] A. Wildani, E. L. Miller, and O. Rodeh, "Hands: A heuristically arranged non-backup inline deduplication system," in Data Engineering (ICDE), 2013 IEEE 29th International Conference on. IEEE, 2013, pp. 446–457. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/6544846/

[16] X. Zhang, Z. Tan, and S. Fan, "NSBS: Design of a Network Storage Backup System," World Academy of Science, Engineering and Technology, International Journal of Electrical, Computer, Energetic, Electronic and Communication Engineering, vol. 9, no. 11, pp. 1245–1254, 2015. [Online]. Available: http://www.waset.org/ publications/10002702

[17] "lessfs - data deduplication for less." [Online]. Available: https://sourceforge.net/projects/lessfs/
[18] "Opendedup-SDFS," 2016. [Online]. Available: https://github.com/opendedup/sdfs
[19] C. Ungureanu, B. Atkin, A. Aranya, S. Gokhale, S. Rago, G. Calkowski, C. Dubnicki, and A. Bohra, "HydraFS: A High-Throughput File System for the HYDRAstor Content-Addressable Storage System." in FAST, vol. 10, 2010, pp. 225–239. [Online]. Available: https://www.usenix.org/legacy/events/fast10/tech/full papers/ungureanu.pdf
[20] S. Quinlan and S. Dorward, "Venti: A New Approach to Archival Storage." in FAST, vol. 2, 2002, pp.89–101. [Online]. Available: https://www.usenix.org/legacy/events/fast02/quinlan/quinlan html/
[21] "Tar - GNU Project - Free Software Foundation." [Online]. Available: https://www.gnu.org/software/tar/
[22] R. Jain, The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling. Wiley, 1990,google-Books-ID: eOR0kJjgMqkC.
[23] "Best open source network data backup and recovery software." [Online]. Available: https://www.baculasystems.com/
[24] "Data Recovery Customers," 2017. [Online]. Available: https://www.baculasystems.com/customers/success-stories

AUTHORS

Heitor Faria was entitled with the Alien of extraordinary ability visa by the US Government for his work in Information Technology and Open Source Software. Master degree at Brasília National University (UNB). "Bacula: Open Source Backup Software" (English & Portuguese) and "Open Licenses & Fundamental Rights" books author (Poruguese). Bacula, Alfresco and Wordpress Training instructor at Udemy, with more than 1200 students in 47 countries. Works as a System Analyst on a brazilian governmental company called SERPRO and for Neocode Software (Canada). Law Graduated. IT Service Manager and Project Management extension degrees. Bacula brazilian community founder. Has plenty of experience as server/backup systems administrator (Windows, Linux, Netware, directory services) and as IT / Project manager. Speaker at several international open source software events. ITIL-F, EMC 05-001, TOEFL and LPIC-III certificated professional.

Jacir Luiz Bordim received the B.S. degree from Passo Fundo University, Brazil in 1994, M.E. from Nagoya Institute of Technology, Japan, in 2000 and PhD from Japan Advanced Institute of Science and Technology in 2003. During 2003-2005, he worked in the development adaptive medium access control for smart antenna technology at ATR - Adaptive Communications Research Labs, Japan. Since March 2005, Dr. Bordim is an Associated Professor at the Department of Computer Science at University of Brasilia, Brazil. His interests include graph theory, distributed algorithms, routing and medium access control protocols, energy aware computing, network design and analysis.
Dr. Bordim has published in extensively in leading journals and international conference proceedings. He is serving on the program committees of international conferences and on editorial boards such as the International Journal of Foundations of Computer Science, the IEICE Transactions on Information and Systems, the International Symposium on Parallel and Distributed Processing and Applications, the International Workshop on Broadband Wireless Access for ubiquitous Networking, Workshop on Advances in Parallel and Distributed Computational Models, among others. Dr. Bordim received the "Outstanding Achievement Award" for his work during the doctorate course.

Priscila Solis Barreto is professsor at the Computer Science Department, University of Brasília. She has a PhD in Electrical Engineering at University of Brasília, Department of Electrical Engineering and master in Computer Science at the Federal University of Goiás, School of Electrical Engineering. Dr. Solis Barreto worked for several years as a Computer Science Technician and Consultant, developing projects and software for different entities, such as Terra Networks, DR Sistemas and CIAT-BID. She was also a professor at the Catholic University of Goiás for several years. Currently, she continues her research at the COMNET Lab. Her main research interests are multimedia traffic models, network performance, network planning and traffic engineering, cloud computing and sensor networks. Priscila is also interested in the development of software tools for simulation.