White Paper

Hardware based GZIP Compression, Benefits and Applications

Tony Summers
Senior Apps Engineer
AHA Products Group, Comtech EF Data Corporation
June 2008

Abstract
Transferring data packets over a network is more efficient when lossless data compression is applied before transmitting. Storing that same data requires less space if compression is applied prior to writing the data to the storage media. Higher data rates and rapidly expanding storage requirements drive the need for better compression solutions. Power conservation strategies also play an important role today in developing data management systems. Technology advances are enabling products to meet these developing and growing needs.

Introduction
Hardware based GZIP compression offloads lossless data compression and frees up the system's Central Processing Unit (CPU). GZIP is a file format standard where the underlying compression algorithm is called Deflate. Deflate compressed blocks are wrapped with a header and footer to become GZIP files. Deflate is an open source compression algorithm and is widely available as a software tool. No license for use of the software is required. When a general purpose CPU performs GZIP compression, typically either the compression performance is scaled back to maximize data throughput speeds or it runs slow. An efficient solution to this is to offload the compression task to a hardware based GZIP co-processor. The co-processor takes in uncompressed input data, compresses it, and outputs the data in compressed form. The hardware co-processor does many tasks in parallel eliminating the multi-pass and iterative nature that is typical of a general purpose CPU executing the Deflate algorithm. This allows the co-processor to operate at much higher data rates.

Compression in storage appliances or in network links must be lossless. The data is typically protected by CRC generation and checking to ensure that after decompression the original data is one hundred percent restored and accurate, and free from corruption or additional artifacts. After a data packet or file is compressed, CRC generated and attached, it may be encrypted for security purposes. If encryption is to be applied to data, it is important that it be applied after compression since encrypted data is highly random and does not compress.

Compression Performance
Compression performance is measured by two metrics, size reduction and data throughput. Size reduction is usually reported as a ratio of the uncompressed original size divided by the compressed size. Data throughput is measured in Bytes per second as measured on the uncompressed side of the co-processor device.

- HTML FILES COMPRESS ON AVERAGE 4.4 TO 1 (Table 1)
- DATA THROUGHPUT IS 5 GIGABITS PER SECOND WITH THE AHA363-PCIE BOARD

Data complexity has no affect on data throughput. Easy to compress data files that compress with a high ratio, pass through the co-processor at the same high data rate as very complex data, which achieves lower compression ratios. Table 1 is a compression ratio performance comparison of AHA's GZIP co-processor results compared to other industry standard algorithms and with standard file sets determined by two prominent universities. These are the Calgary Corpus and Canterbury Corpus. The HTML file set is from a collection of internet dynamic content.

TABLE 1: Compression Ratio Performance

STANDARD COLLECTION OF FILES	**GZIP CO-PROCESSOR**	**LZS**	**ALDC**
CALGARY CORPUS	2.73	2.24	2.10
CANTERBURY CORPUS	3.6	2.75	2.68
HTML file set	4.4	N/A	2.65

Power consumption should be considered when looking for a GZIP compression solution. Comparing the AHA GZIP compression chip to a quad core CPU dedicated to running GZIP compression software shows that the AHA device requires about 8x less power. In a data center the potential for power savings are great. Offloading compression to a co-processor instead of using the CPU saves a significant amount of power. Transmitting smaller compressed packets results in additional power savings. Storing compressed data to the media also saves power.

THREE TO ONE DATA COMPRESSION RESULTS IN ONE THIRD LESS DATA PASSING THROUGH THE DATA CENTER, AND STORING DATA REQUIRES ONE THIRD THE SPACE ALLOCATION.

OFFLOADING GZIP COMPRESSION TO A CO-PROCESSOR RESULTS IN 8X LESS POWER CONSUMPTION.

Applications and Solutions

Many network and storage appliance companies would prefer to implement a plug-in GZIP compression board instead of developing their system around a GZIP compression-decompression ASIC device. AHA offers a PCI Express board with drivers and driver development support that simplifies the design-in process allowing companies to take advantage of this new technology with less design in effort. This GZIP board (AHA363-PCIe) achieves 5 Gigabits per second data throughput using two AHA co-processor devices. The board is capable of full duplex operation, compressing and decompressing multiple streams or files.

Reducing retrieval time of dynamic content when accessing a web site is one problem that network appliances can improve on by implementing a compression accelerator. This is also referred to as application acceleration and is important when retrieving dynamic content such as financial data. Implementing hardware based GZIP compression is an ideal solution since it offloads the compression function to a co-processor and keeps the CPU free for other tasks. Another reason for choosing GZIP (Deflate) as the algorithm is that most web browsers have the decompressor built in and can automatically uncompress these packets as they are received. This works well since the decompression process for GZIP is not nearly as CPU resource intensive as the compression process. Figure 1 shows compression accelerator boards installed in different network appliances. Some of the installations use hardware compression to speed up network traffic. Others use hardware compression to reduce the space required on the storage media.

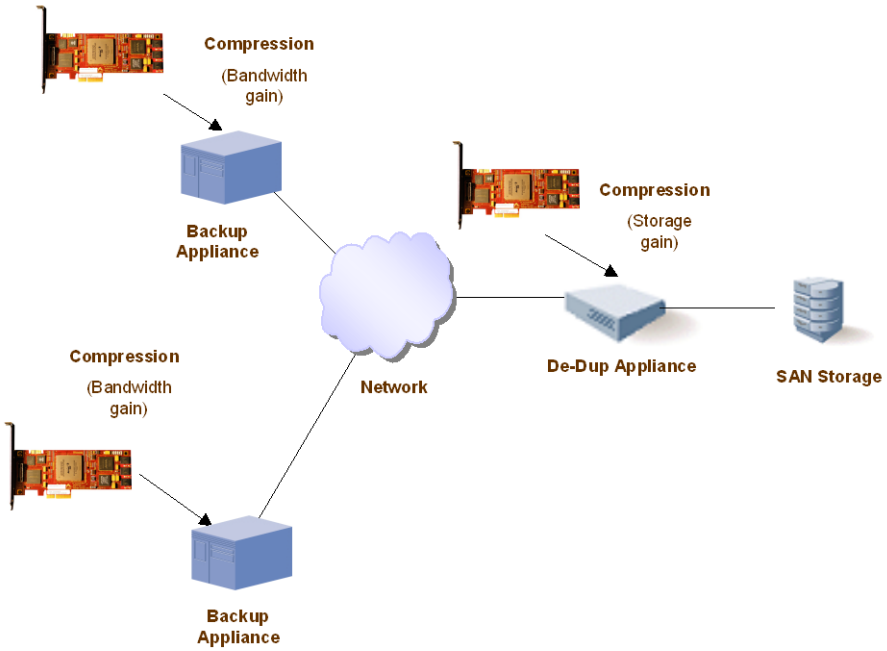

Figure 1: Compression Accelerator boards in Network Appliances

Storage appliances such as Virtual Tape Libraries and Wide Area Network (WAN) Acceleration Appliances can benefit from a PCI Express based GZIP compression board by compressing the data content before writing to the media. For this application high compression ratio is the most important requirement followed by high data rate. In the past LZS and ALDC have been popular compression solutions, but both algorithms fall short compared to AHA's latest product offerings that implement GZIP (or Deflate).

Conclusion

The high growth of corporate data bases, back-up recovery, and electronic mail as well as long term data retention requirements are driving the need for ever growing data storage capacity. Data centers are being built near hydroelectric power plants due to their huge demand for power. Lossless data compression is now a necessity in network or storage system developments. GZIP (Deflate) is the most popular and best performing lossless data compression standard. Offloading GZIP compression to a co-processor greatly improves data throughput efficiency, higher compression ratios, and power conservation. Figure 2 shows a compression board installed in a storage appliance to reduce the space required to store the data.

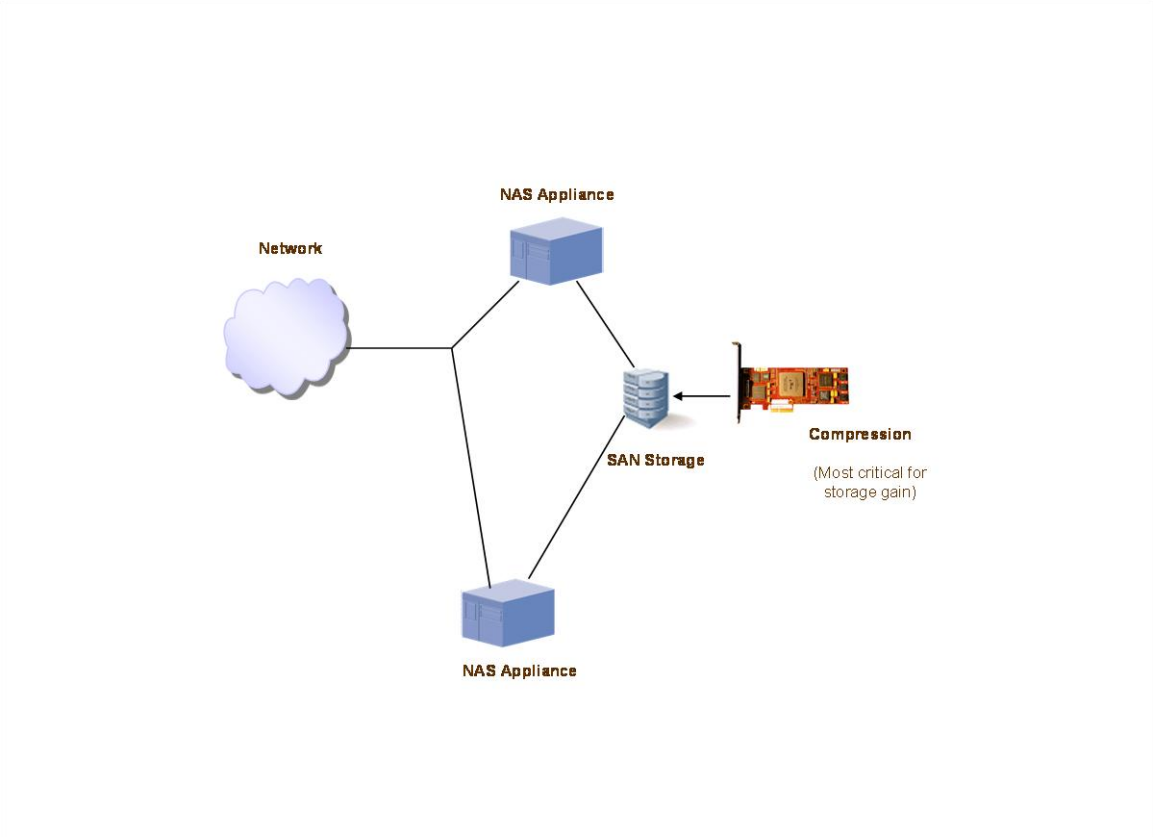

Figure 2: Acceleration board in a Storage Attached Network (SAN) Appliance

Company background

The AHA Products Group (AHA) of Comtech EF Data Corporation develops and markets superior integrated circuits, boards, and intellectual property cores for improving the efficiency of communications systems everywhere. AHA has been setting the standard in Forward Error Correction and Lossless Data Compression for many years and provides flexible and cost effective solutions for today's growing bandwidth and reliability challenges. Comtech EF Data is a wholly owned subsidiary of Comtech Telecommunications Corporation (NASDAQ" CMTL). For more information, visit: www.aha.com.