

Validating the Effects of Dedupe/Compression on Storage Performance

Extending the business case for solid state storage

A key feature of today's all-flash storage array implementations is the use of data deduplication and compression. Deduplication and compression are technologies that dramatically improve IT economics by minimizing storage requirements, backup windows, and network bandwidth consumption (not so much – Reads and Writes to storage still occur. Backup size is smaller in most cases, but this is not a major benefit. Instead, we might want to discuss things like dramatic savings in cooling, footprint and reduced need for compute infrastructure). These technologies are particularly attractive in the solid state storage space to enhance SSS financial viability and extend their use to a greater number of application workloads beyond the Tier-1 applications where they are most commonly initially deployed.

SSS operates differently than previous generations of spinning media arrays. With flash memory, the rotational and stroke latencies inherent in disk-based access are gone, which means access to any data on an array is just as fast as any other. Additionally, faster processors and sophisticated data reduction algorithms enable these arrays to find duplicate data patterns and eliminate them by storing them only once. Some arrays even have the ability to find patterns when they are skewed, or offset. For example, if someone adds data to the beginning of a file, these arrays can find the offset duplicate pattern and refer to it rather than storing it again.

As with deduplication, many modern all-flash arrays store repeating characters as metadata, referring to the repeating character and how many times it repeats, instead of storing the repeating pattern. Additionally, some modern arrays can apply their processing to compression, which is data reduction involving encoding information using fewer bits than the original representation.

But you can't treat dedupe/compression like a checkbox. Implementations vary as well as the effect on different workloads.

“Accurately measuring the performance of deduplication and compression depends on generating data content patterns sufficient to stress a storage array.”

Leah Schoeb
EVALUATOR GROUP

Load Dynamix is uniquely positioned to properly test and validate modern storage technologies like deduplication and compression with high accuracy, tremendous load and extremely realistic workload patterns. We use a patent pending method of generating and verifying streams that represent repeatable random and repeatable compressible content. These algorithms can generate exabytes of unique data, enabling testers to specify the wide range of patterns and repetition required to properly exercise data reduction technologies, thereby helping to test arrays from any vendor that implements the features. We generate a combination of completely unique patterns and repeating patterns. Any repeating pattern should reduce down to one deduplicated entry. We do the same with compressible data, either unique or repeating. And we combine those data types to generate a data set and the data streams required to properly test an array. Load Dynamix uses multiple Write statements, for the selected protocol, with the capability of selecting the appropriate data type and the number of times repeated to lay down a dataset via a data stream that emulates real world traffic. The storage array then executes the deduplication and compression based on those patterns, stores as metadata, and we measure the results on metrics like IOPS (see figure one on the next page).

Here's an example that illustrates our methodology, using a 100% sequential write workload, in a block environment:

Assume this Logical Unit (LU) and its Logical Block Addresses (LBAs):



If your goal is to test data content that has a 2:1 deduplication ratio, Load Dynamix would, starting at the first address, write a random, unique block (B) and at the second address, a deduplicable record (D), and so on.



After deduplication this results in 4 unique blocks, plus 1 deduplicable block. So the ratio would be 8:(4+1), or 8:5, not quite 2:1. If your workload truly only consisted of these 8 block, 8:5 would be your ratio. But in a more realistic environment, your LUs would be much larger and your workload would generate thousands of writes. With 1000 writes, your ratio would look like 1000:(500+1), or nearly perfectly 2:1.

The same logic would apply in a random write environment, even if the performance may be different.




Figure 1: Example test, comparing vendor results at varying dedupe and R/W ratios.

Load Dynamix Deduplication Test / Validation Example

In the graph above, we demonstrate that different arrays will perform differently using different amounts of repeating data. In this example, vendor B exhibits higher IOPS when data is 20% or 50% reducible, with any read/write ratios, but especially with higher reads. Vendor A shines when data is highly reducible and read/write ratios are in the 50% range (or greater).