



Data Reduction Engine

MangoBoost's patent-pending Data Reduction Engine (MB-DR) offers inline data deduplication and compression at 2GB/s, scalable to many SSDs

Key Features

- **Patent-pending** data reduction engine with **integrated deduplication and compression FPGA offload**
- Allows the **integration of custom hash/compression IPs** using standard **AXI interfaces and with configurable number of units.**
- Provided examples with FIPS 180-2 compliant signature generation using **SHA-256**, and **LZ4** compression
- **512K IOPS** data reduction throughput with 4KB random reads and writes
- **Lightweight software support** for easy integration (< 100 lines of code for integration with Red Hat Virtual Data Optimizer (VDO))
- **55% space savings** shown with SPECstorage Solution 2020 Benchmark EDA dataset
- Allows **configurable data chunk size** (4KB-64KB)
- Supports AMD **Alveo-series FPGAs**
- Code in C-based HLS (AMD Vivado)
- SystemC version available soon (tested using Siemens Catapult HLS) for portability to other FPGA vendors

Application Use Cases

- Large scale data store in big data workloads, like AI, genomics, health, etc
- Storage server ??? AI Integration with Redhat Virtual Data Optimizer (VDO)
- Integration with Linux Device Mapper modules (dm-dedup and dm-verity)
- Accelerated signature generation for data integrity verification

Case Study: SPECstorage 2020 & VDO acceleration

- **Setup:** MB-DR on AMD U280 FPGA with four NVMe SSDs (total 6 TB), in a host Intel Xeon Gold 6338 2.0 GHz CPUs
- Installed NFS server and client over a single 100-GbE network
- MB-DR runs at 250MHz, utilizing 255K LUTs, 416K Regs, 758 BRAMs, 48 URAMs
- Achieved 20% higher SPECstorage benchmark score versus software baseline with 35% CPU savings
- 55% storage savings with electronic design automation (EDA) dataset

