

MongoDB Acceleration with MonetX

FPGA based MongoDB acceleration
with zero code change

INTRODUCTION

Databases provide a wealth of functionality to a wide range of applications. Yet, there are tasks for which they are less than optimal, for instance when processing becomes more complex or the data is less structured. As data is exploding exponentially only CPU based systems no longer provide real-time insights to businesses in a cost-effective way. At Grovf we designed a Monet – A FPGA based smart memory extension for near memory data processing. Monet implemented on top of Xilinx's Alveo U50 acceleration card and once plugged into server's PCIe bus acts as a standard RAM memory for the Linux operating system with in-memory compute API capability.

KEY BENEFITS

- 10GB/s, 2TB RAM memory on a single PCIe slot
- Network accessible memory
- In-memory data computing
- RAM Memory extension without increasing number of servers' sockets
- Linux compatible


SOLUTION BRIEF



- 10GB/s, 2TB RAM memory on a single PCIe slot
- In-memory data computing
- Linux compatible

SOLUTION OVERVIEW


MonetX acts as a standard RAM memory once connected to the server's PCIe slot. With in-memory computing capability it provides simple API to host layer for easy utilization of the functions. Based on the Monet smart memory extension MongoDB has been accelerated 3.5X for all stages of data aggregation.



FPGA based MongoDB acceleration with zero code change


SOLUTION DETAILS

MongoDB acceleration is based on Grovf's MonetX acceleration platform, which is an FPGA based smart memory extension for near memory data processing. The operating system recognizes MonetX as a standard memory extension which also provides high-performance computing cores API for the host layer. Data can be stored to MonetX memory extension, just like into any other memory connected to the server. Applications that can initiate different processing on the data stored into the MonetX platform directly running on FPGA, where data resides also. MonetX supports many high-performance computing cores such as Regular Expression processing, Search/Sort processing, Data compression/decompression, Statistical Data processing algorithms, etc. MongoDB performance has been boosted 3.5X only using the MonetX acceleration platform as a high bandwidth memory extension for standard server architecture without using any build-in high-performance computing cores in the FPGA. This leads to zero code change in the application(MongoDB) side and provides 3.5X acceleration. More acceleration for the MongoDB and any other application can be achieved using build-in accelerated computing cores in the FPGA residing near memory.



RESULTS

When benchmarking the MongoDB without MonetX smart memory extension card, the performance of the database starts decreasing. Drastic decrease happens when any stage of the MongoDB function uses more than 100MB of memory. With MonetX MongoDB persists it's computational power even when processing occupies more than 100MB of RAM for each aggregation stage.



TAKE THE NEXT STEP

Learn more about Xilinx [Alveo accelerator cards](#)

Learn more about [Grovf Inc.](#)

Reach out to [Grovf Inc. sales](#) – artavazd.rk@grovf.com, khachik.ss@grovf.com