

Query Processing IP Built on NoLoad® Computational Storage Framework

Eideticom's Query Processing Unit (QPU) targets database users or anyone streaming data (such as network packets) who need query, analytics or format conversion done in hardware with low latency. Tasks can be parallelized to meet any bandwidth needs, plus mixed with other NoLoad® functions like compression.

Built on the NoLoad® Framework

As part of Eideticom's NoLoad computational storage framework, the QPU can be pipelined with other Computational Storage IP from Eideticom such as Compression, Decompression Erasure Coding, and Deduplication.

Software-Driven Design

Relieving the burden of hardware engineers being involved in specifying the QPU parameters, users define the functions using an on-chip processor. This gives both high-level tool ease-of-use and offload from the main host CPU.

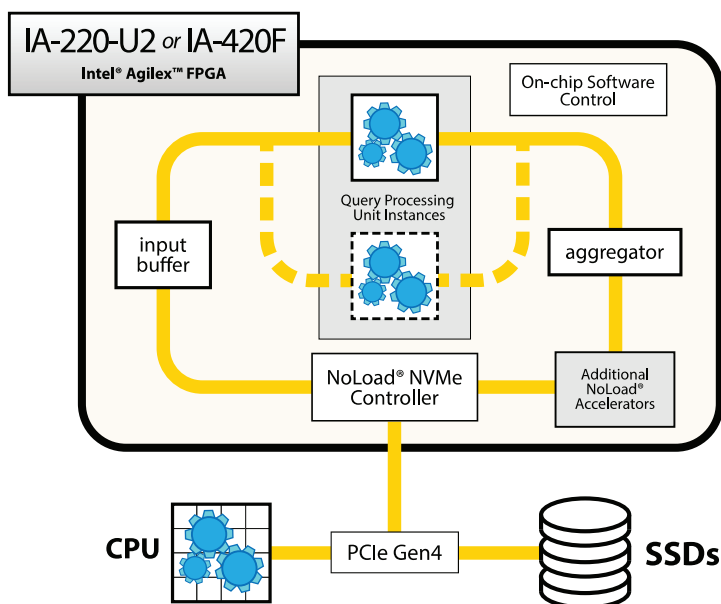
The Query Processing Unit includes features for format conversion (text to/from binary) or standards-driven database functions. Users design their mix of functions in software—no hardware-based tools are necessary. The QPU will support native acceleration for a range of database tools as these packages move to support standards-based computational storage.

key features

Software-defined search and filter parameters

Efficient reformatting of text/binary data

Filter, pattern-match, and analytics for streaming data



How QPU Works Inside of NoLoad, Deployed on BittWare's IA-220-U2 Module or IA-420F Card

Features

- Data from user space applications is stored using many different formats
- NoLoad Query Processing Unit accelerates data queries
 - Configurable text and binary input and output formats
 - Configurable search and filter parameters
- QPU can be paired with NoLoad compression and decompression engines
- libnoload user space API provides access to accelerators
 - Uses NVMe driver to achieve low latency high throughput data transfers to the QPU
 - Initially targeting Fintech packet analytics applications
- Path to NVMe TP4091 for broad application deployment

Query Processing Unit

IP Core

Software Defined + Scalable for Your Bandwidth Requirements


The QPU is defined in software that runs on the FPGA (using soft or hard processor), eliminating the need for low-level configuration engineering resources. The Query Processing Unit is modular, allowing for one or more units to be placed to meet a certain bandwidth requirements. Unit instances can cooperate, for example with a distributed file conversion over eight QPUs where data spanning two units needs to be coordinated.

Built on the NoLoad® Computational Storage Framework from Eideticom

The Query Processing Unit is a component of the NoLoad framework. The components such as Compression in orange are where users build their particular application using a software-defined approach.

Components like Compression can be added to the QPU to, for example, compress filtered data before moving to SSD storage.

All the accelerator functions shown are implemented in FPGA hardware, allowing for **high bandwidth, low latency and CPU offload**.



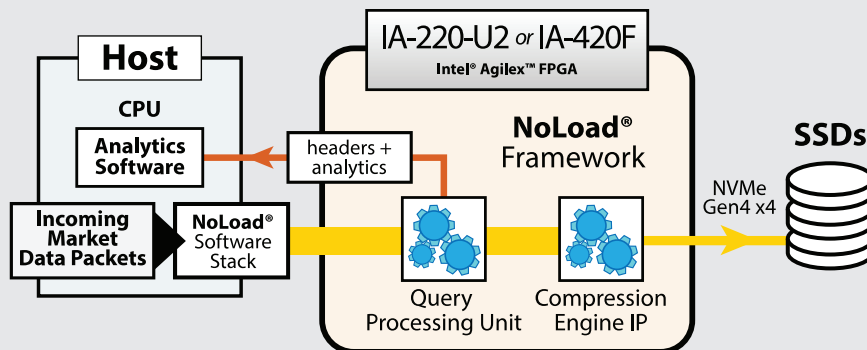
Use Case: Capture + Analytics Engine for Fintech

This real-world example uses the Query Processing Unit as a packet processing machine, plus the Compression Engine (another NoLoad® IP core). Packets are compressed and written to an SSD array using peer-to-peer transfers, while the QPU also pulls off header (network tuple) data plus some analytics such as packet count per time period. Analytics are sent to the host as CSV data.

Performance Advantage

Compared to a multi-threaded Intel Xeon CPU, the database Query Processing Unit performs as shown below.

Avg. Packet Size	CPU	1× QPU	2× QPU	4× QPU
256B	0.2 GB/s	1.8 GB/s	3.6 GB/s	7.2 GB/s
1024B	0.7 GB/s	2.0 GB/s	4.0 GB/s	7.2 GB/s
4096B	1.9 GB/s	2.0 GB/s	4.0 GB/s	8.0 GB/s
9216B	2.7 GB/s	2.0 GB/s	4.0 GB/s	8.0 GB/s



Packets from the host are moved over PCIe to the BittWare IA-220-U2 or IA-420F for FPGA processing using the NoLoad® framework and IP

Query Processing Unit

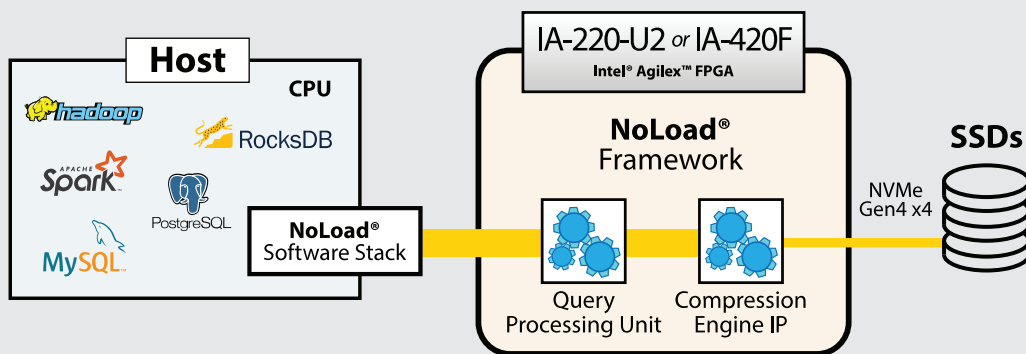
IP Core

Use Case: Database Query Acceleration

Accelerate Database Queries and More

In a database acceleration configuration, the Query Processing Unit can perform a range of functions from CPU offload to data type format conversions.

- Implements **CSV/JSON/Parquet parsing** and query execution
- **70-80% Improved CPU Offload** with **5-10x Increased Performance**
- **Deployment flexibility** in compute, storage and cloud
- NVMe driver achieves **low latency high throughput** data transfers
- **NoLoad Query Processing Unit** can be paired with NoLoad Compression and Decompression Engines



The NoLoad™ Software Stack brings DB data to/from the accelerator and storage

Query Processing Unit Feature List

- Provides SQL-like semantics to query data
 - `SELECT columns FROM source WHERE conditions;`
- Input Data
 - Row-based text formats (CSV, JSON, etc.); CSV is the most popular data format for data storage!
 - Row-based binary data (C structures, Avro, etc.)
- Input SELECT Engine
 - Easily configurable to define input data format
 - Generates internal binary tokenized data
 - Filters columns completely unused in filtering or output
- WHERE Engine
 - Apply filters to tokenized data to extract only necessary data rows

- Output SELECT Engine
 - Easily configurable to define output data format
 - Generates output data format
 - Filters columns and returns only requested columns
- Output Data
 - Row-based text formats (CSV, JSON, etc.)
 - Row-based binary data (C structures, Avro, etc.)

Compatible FPGA Cards

- [IA-220-U2](#)
- [IA-420F](#)

Looking for a different card? Ask us about other compatible card options

To learn more, visit www.BittWare.com

Rev 2022.5.26 | May 2022

© BittWare 2022

NoLoad is a registered trademark of Eideticom. All other products are the trademarks or registered trademarks of their respective holders.

BittWare
a **molex** company