



Low-Latency RoCE v2 at 100Gbps

The GROVF RDMA IP core and host drivers provide RDMA over Converged Ethernet (RoCE v2) system implementation and integration with standard Verbs API. The RDMA IP is delivered with a reference design that includes the IP subsystem itself, the 100G MAC IP subsystem, the DMA subsystem, host drivers, and example application on software. The system drivers are integrated with OFED standard Verbs API and are compatible with well-known RNIC cards and software. The IP core also provides a low-latency FPGA implementation of RoCE v2 at 100 Gbps throughput.

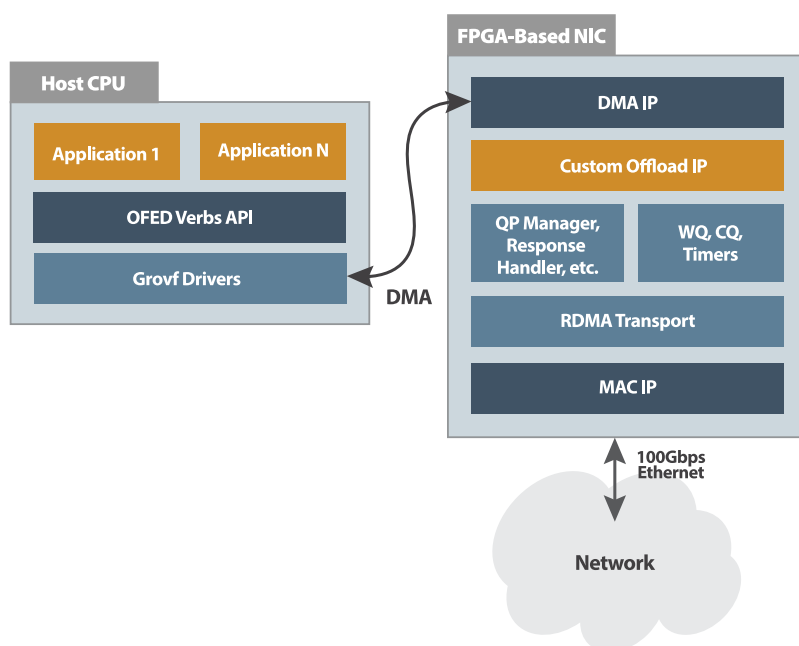
Enables RNIC
use cases with
FPGA-based
SmartNIC

key features

Compatible with
RNIC and soft
RoCE v2

100Gb/s
throughput,
2μs latency

1023 or more
configurable
RDMA queue pairs



Features

- Hardware operated RC, XRC, RD, UC, UD services
- Incoming and outgoing SEND, RDMA READ, RDMA WRITE
- Memory protection domains implemented in FPGA and ECN
- 3rd party MAC and DMA IPs
- Standard Verbs API on Host Machine
- Dynamic configuration using Verbs API
- Hardware retransmission and reordering
- Customizable IP

RDMA

IP Core for RoCE v2

Product Operation

The solution is a soft IP implementing RDMA over Converged Ethernet protocol. It consists of FPGA IP integrated with MAC and DMA, plus the host CPU drivers. The IP is compatible with BittWare's IA-420f FPGA cards* featuring Intel Agilex 7 and with XUP-VV8 and XUP-P3R FPGA cards featuring AMD UltraScale+. The solution complies with Channel Adapter and RoCE v2 requirements as stated in the IB specification. The diagram on page 1 shows a simplistic architectural overview of the system. The data plane and reliable communication is hardware offloaded and the implementation does not include CPU cores in the FPGA.

Detailed Feature List

- Fully compatible with known RNIC products and soft RoCE implementations (RoCE v2)
- Under 2.0 μ s software to software latency (roundtrip) and under 300 ns hardware to hardware latency (roundtrip)
- 100 Gb/s throughput
 - Configurable RDMA queue pairs
 - 1023 or more
- Hardware retransmission management
- Memory protection domains implemented in FPGA
- Congestion control using ECN, PFC
- Can work with 3rd party MAC and DMA IPs
- Dynamic configuration using Verbs API
- Standard Verbs API on host machine user / kernel space
- Hardware implemented Reliable Connection (RC), Extended Reliable Connection (XRC), Reliable Datagram (RD), Unreliable Connection (UC), and Unreliable Datagram (UD)
- Incoming and outgoing SEND, RDMA READ, RDMA WRITE. RDMA Atomic is not supported.

Reference Designs

The reference example consists of three parts:

- Encrypted FPGA IPs with reference design which implement RDMA protocol
- Software drivers which provide standard Verbs API support for the FPGA based RDMA adapter
- Example application build on top of the Verbs API demonstrating ping-pong test results: latency and bandwidth

Sample Implementation Results

Device	LUTs	On-Chip Memory
UltraScale+ VU9P	170K	6Mb
Agilex 7 AGF014	170K	6Mb

Compatible FPGA Cards

- [IA-840f*](#)
- [IA-420f](#)
- [XUP-VV8](#)
- [XUP-P3R](#)

Looking for a different card? Ask us about other compatible

* IA-840f support available soon.

To learn more, visit www.BittWare.com

Rev 2024.02.19 | February 2024

© BittWare 2024

UltraScale+ is a registered trademark of AMD corp. All other products are the trademarks or registered trademarks of their respective holders.

BittWare
a **molex** company