



The Open Fabrics Verbs Working Group

Pavel Shamis and Liran Liss

Introduction

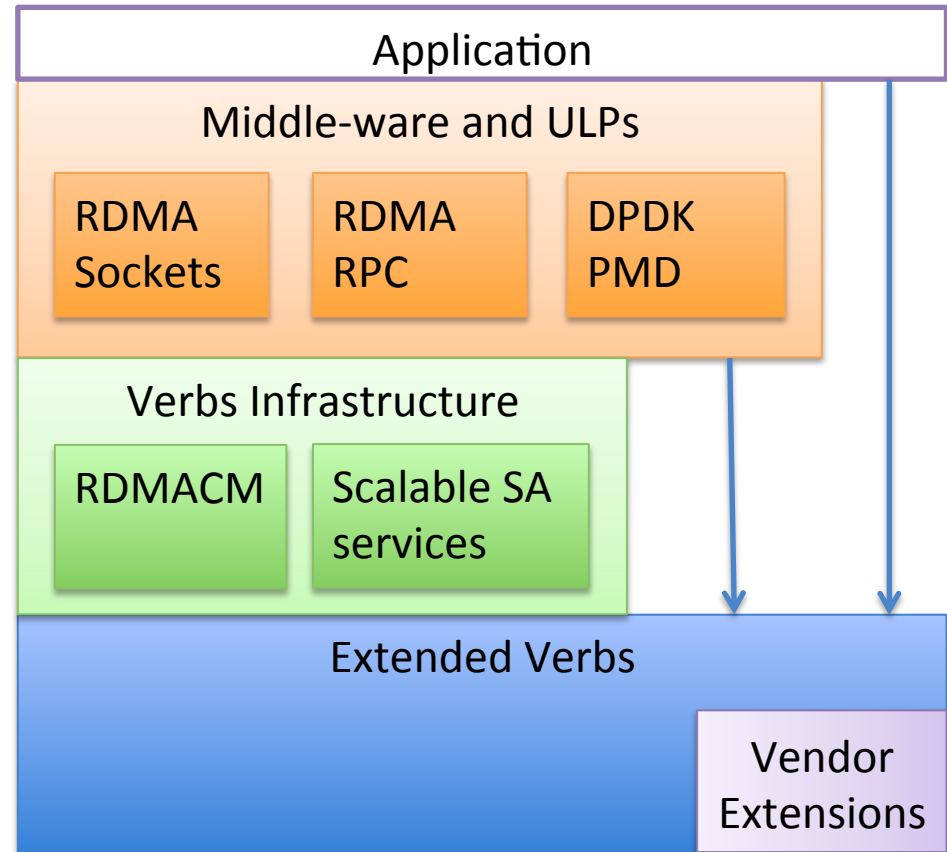
- Verbs is a high-performance mature and robust interface
 - Widely adopted
 - Stable
 - Scalable
- Open-source development
 - Peer-reviewed code patches
 - Peer tested
- Development rate and novelty constantly increasing
 - New applications and features
 - New systems and architectures
- The Open Fabrics Verbs Working Group (OFV-WG)
 - Enhancing the verbs interface to meet the needs of future systems and applications

OFV Mission

- Enhance the Verbs interface to meet the needs of future systems and applications
 - HPC, storage, cloud applications and more
 - CPUs, GPUs, and other compute elements
 - RDMA and other interconnect accelerated capabilities
- Guide the development of the Verbs eco-system
 - Forum to discuss approach for new features
 - Raise community awareness
 - Partner with collaborators for introducing new features

The Verbs Eco-System

- Extended Verbs
 - Enable new features
- Infrastructure
 - Common services
 - Common abstractions
- Middleware and ULPs
 - Application-facing APIs



The Verbs Eco-System: Extended Verbs

- HPC optimizations
 - Accelerated Verbs
- Raw Ethernet support
 - Rx flow steering
 - Tx flow anti-spoofing
 - Stateless offloads
 - VLAN stripping
 - RSS, TSS
 - Tunneling (VXLAN, NVGRE)
- Memory
 - Indirect Memory Regions
- Signature operations
- Extended atomics
- Name space operations
 - QPNs, CQNs, MRs, etc.
- QP control
 - Suspend/resume
- Time stamp operations
 - Read HCA clock
 - Time stamp completions

The Verbs Eco-System: Verbs Infrastructure

- RDMACM improvements
 - E.g., APM support
- Bulk address resolution for HPC
- Kernel-managed user-space QPs
 - Connection management
 - fork() support
- SoftRoCE
- Multi-rail bonding abstraction
- On-Demand-Paging
- Container support
- SELinux support

The Verbs Eco-System: Verbs Middleware



- Group communication
- Transaction-based IO
 - Efficient RPC
- Active messages
- RDMA sockets
- User-space TCP/IP
- Packet processing
- Java bindings
- RDMA accelerators
 - E.g., HPC, Hadoop, Ceph

Discuss New Approaches

- Solicit feedback from a larger community and users
 - Not all of our contributors track kernel mailing lists
 - Storage, Big-Data, HPC, etc.
- Discuss concepts before writing code or detailed [RFC] patches
- Converge faster by interactive feedback
 - Faster acceptance
- Tackle hard-to-crack concepts by focused discussions
 - Weekly increments

Community Awareness

- Raise the need for important features
 - Introduce new use-cases
 - Describe feature requests
 - Aid in prioritizing the focus of the development community
- Point out pain points and urgent issues
 - E.g., connection rate, “debugability”

Collaboration

- Raise issues that you would like to work on
- See if other community members are currently tackling the same issues
- Establish collaboration to speed up development and increase efficiency
- Cross-community collaboration
 - Discuss challenges with broader community

Summary

- The pace of Verbs development is increasing
- The user base of the Verbs Eco-system is expanding
- The OFV WG will facilitate
 - Agreement on approach for new fronts
 - Faster acceptance and convergence
 - Prioritize and focus OFA development efforts
 - Encourage collaboration
 - Broaden the use of the interface
 - Increase adoption of RDMA technology



Thank You



OPENFABRICS
ALLIANCE