OPENFABRICS
ALLIANCE

12th ANNUAL WORKSHOP 2016

COMPLIANCE AND INTEROPERABILITY

Paul Bowden, Bob Russell, Paul Grun

April 5, 2016

THE OPENFABRICS ALLIANCE PERSPECTIVE

The OFA strives to deliver services that bring value to alliance members

- **One way it does that is through an Interoperability program**
- **Two key goals of the Interop program**
  - Grow adoption of OpenFabrics Software (OFS) by building confidence that a valid ecosystem exists
  - Support vendors that deploy products incorporating OpenFabrics Software
- **With the emergence of the OpenFabrics Interfaces project, this may need a re-think**

Expand the existing interop program to add more value to the alliance community

THE ALLIANCE COMMUNITY

emerging

historical

- Developers of code that directly accesses network services
  - Middleware developers
  - Application coders
- People who sell systems that include networks
  - OEMs, VARs
- People who buy, deploy, use or maintain systems that include networks
- Network hardware & software vendors
  - IHVs, ISVs
  - Open source developers

One of the few places in the world where the interests of vendors and consumers intersect

OPENFABRICS INTERFACES PROJECT - OFI

- **Formed in August 2013 to encourage the development of network APIs that are:**
  - Transport independent, and
  - 'Application Centric' – i.e., highly responsive to the needs of 'apps' that will use them

- **Planned to include a family of APIs, targeted at specific use cases**
  - **libfabric** – a user mode library focused on Distributed and Parallel Computing uses
  - **kfabric** (proposed) – kernel mode functions focused on Data Storage and Data Access
  - future ? – APIs for:
    - Data analysis,
    - Big Data,
    - Cloud computing,
    - Virtualization…

Given this, is there an opportunity for the OFA to deliver more value to the community?

IWG PROGRAM (TODAY)

- **Built mainly around vendor testing**
  - Purpose is to demonstrate interoperability…and a sound ecosystem
  - Vendor products are improved through debug events and testing
  - Deployers of RDMA networks also benefit

- **Consists of -**
  - Debug events – vendors can test and debug in a heterogeneous system setting
  - Logo program – a badge of honor awarded to a device
    - Method: exhaustive testing conducted by a neutral party (UNH-IOL) on a cluster in neutral territory

COMPLIANCE? INTEROPERABILITY?

Interoperability:

- **A device is interoperable if it works "correctly" with other devices – e.g. does this HCA work with other HCAs? Switches?**
- **Interoperability is established through a series of exhaustive tests, usually focusing on a matrix of components**

Compliance:

- **An object is in compliance if it conforms to a set of requirements as measured by some objective criteria**
- **Typically, requirements are conveyed in an industry standard**
- **Compliance is determined by testing against each requirement**

Today, devices are not tested for compliance
Compliance testing is left to the standards bodies e.g. IBTA, IEEE…

INTEROPERABILITY – IB EXAMPLE

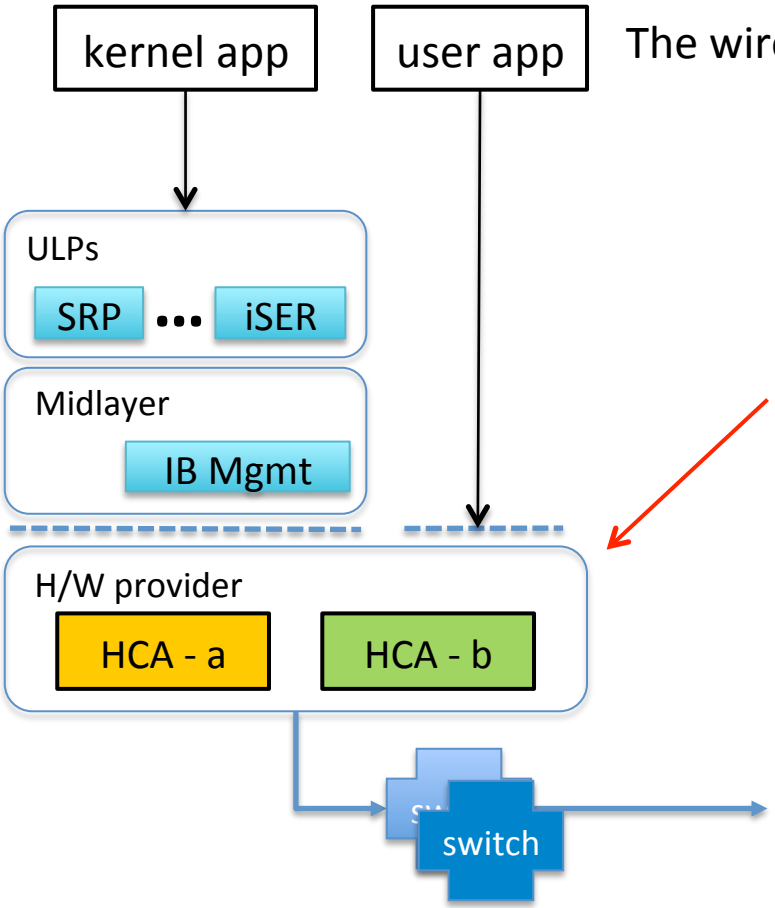

The wire dictates the services available to the user

Multiple vendors for each component type

→ interoperability makes sense

→ compliance testing (does the device conform to the spec?) is left to the spec owner, typically a standards body

Same wire protocol, independent of vendor

WHO BENEFITS FROM CURRENT INTEROP?

users of OFS

creators of OFS

- Developers of code that directly accesses network services
  - Middleware developers
  - Application coders
- People who sell systems that include networks
  - OEMs, VARs
- People who buy, deploy, use or maintain systems that include networks
- Network hardware & software vendors
  - IHVs, ISVs
  - Open source developers

OFI - EXPANDING OPENFABRICS SOFTWARE

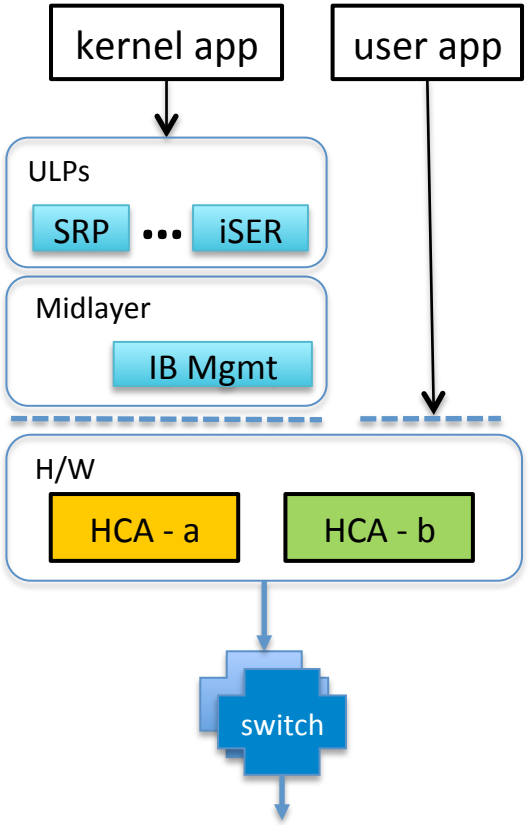

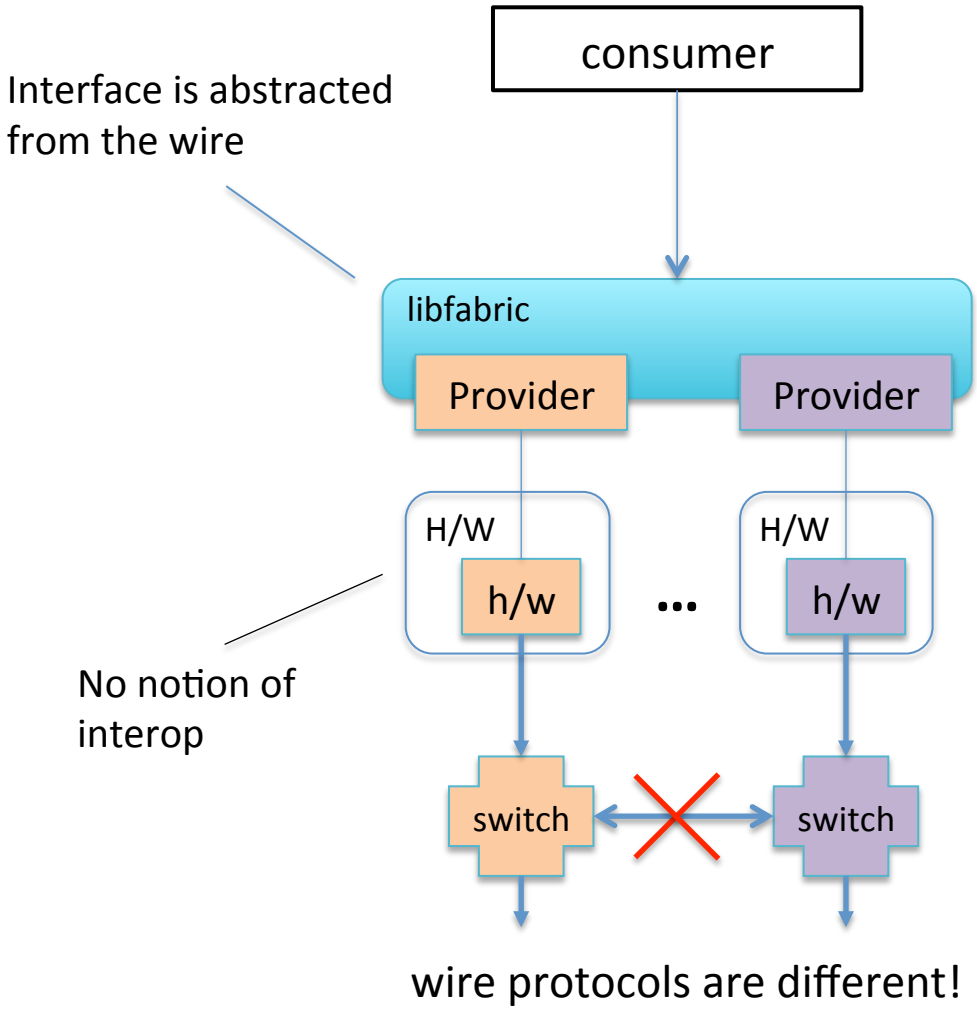

COMPLIANCE

The goal now is to ensure that each vendor ('Provider') implements the API functions correctly, and that services are exported to the consumer correctly.

In other words, ensure that the design complies with the API as defined in the libfabric MAN pages.

Consumer code (middleware, application level) now becomes portable between providers.

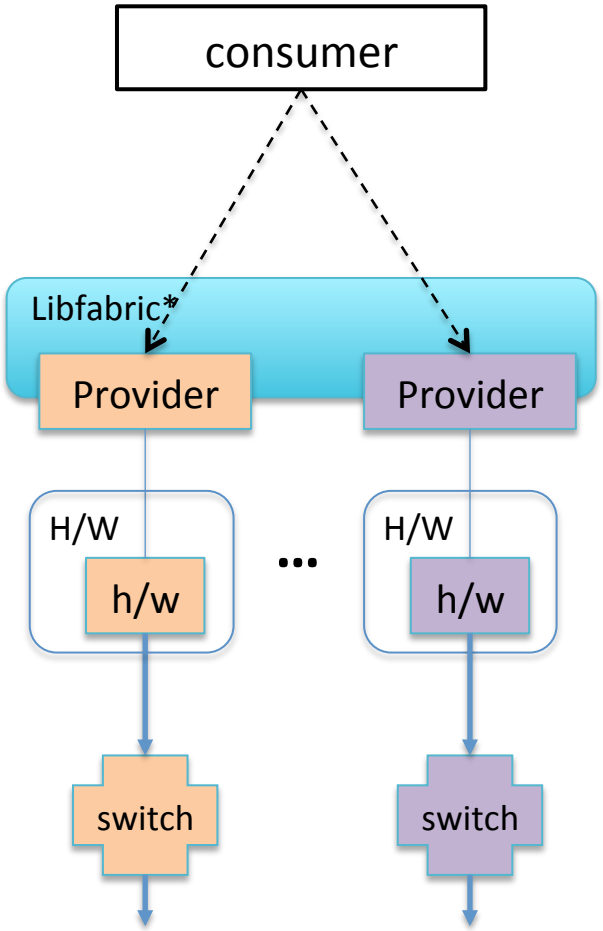

*libfabric – user mode library for distributed and parallel computing

COMPLIANCE OBJECTIVE

Endpoint Types	gni	mxm	psm	psm2	sockets	udp	usnic	verbs
FI_EP_DGRAM	✓	✗	✓	✓	✓	✓	✓	✗
FI_EP_MSG	✗	✗	✗	✗	✓	✗	-	✓
FI_EP_RDM	✓	-	✓	✓	✓	✗	-	-

Primary Capabilities	gni	mxm	psm	psm2	sockets	udp	usnic	verbs
FI_MSG	✓	✗	✓	✓	✓	✓	✓	✓
FI_RMA	✓	✗	✓	✓	✓	✗	✗	-
FI_TAGGED	✓	✓	✓	✓	✓	✗	✗	-
FI_ATOMIC	✓	✗	✓	✓	✓	✗	✗	-
FI_NAMED_RX_CTX	✗	✗	✗	✗	✓	✗	✗	✗
FI_DIRECTED_RECV	✓	✗	✗	✓	✓	✗	✗	✗
FI_READ	✓	✗	✓	✓	✓	✗	✗	✓
FI_WRITE	✓	✗	✓	✓	✓	✗	✗	✓
FI_SEND	✓	✗	✓	✓	✓	✓	-	✓
FI_RECV	✓	✗	✓	✓	✓	✓	-	✓
FI_REMOTE_READ	✓	✗	✓	✓	✓	✗	✗	✓
FI_REMOTE_WRITE	✓	✗	✓	✓	✓	✗	✗	✓

- Multiple providers
- Delivering unique features

Compliance Objective: if a provider supports a feature, is that feature implemented correctly?

THE ALLIANCE COMMUNITY

users of OFS

creators of OFS

- Developers of code that directly accesses network services
  - Middleware developers
  - Application coders
- People who sell systems that include networks
  - OEMs, VARs
- People who buy, deploy, use or maintain systems that include networks
- Network hardware & software vendors
  - IHVs, ISVs
  - Open source developers

Can we expand the interop program so it delivers (more) value up the stack?

DRIVING CONSUMER BENEFIT

Take a look at a typical code development cycle

- **Stage 1** – Architecture
  - Should I deploy OFI?
- **Stage 2** – Decision made, now what?
  - How to design code for OFI?
- **Stage 3** – Code Development
  - Development environment, experimental cluster
- **Stage 4** – Test, validate, deploy
  - Proper operation, optimization
- **Stage 5** – Maintenance, feature development

POSSIBLE OFA VALUE ADDS

Take a look at a typical code development cycle

- **Stage 1** – Architecture
  - Should I deploy OFI?
- **Stage 2** – Decision made, now what?
  - How to design code for OFI?
- **Stage 3** – Code Development
  - Development environment, experimental cluster
- **Stage 4** – Test, validate, deploy
  - Proper operation, optimization
- **Stage 5** – Maintenance, feature development

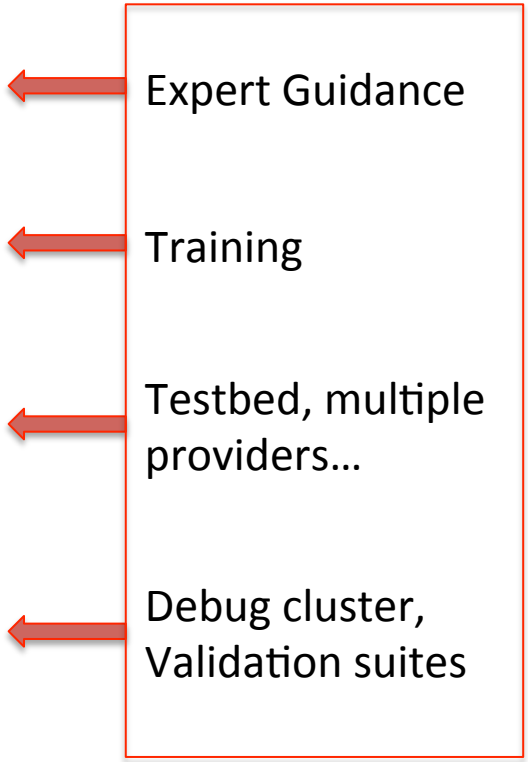

(INFORMAL) PROPOSAL

Augment the existing Interop Program with these elements:

- **Training**
  - Key beneficiary is the consumer community
- **Provider Compliance Validation**
  - Key beneficiary is the provider vendors (with some benefit to the consumer)
- **Support for Consumer Code Development**
  - Assistance and advice in implementation and coding
  - Emphasize agile development techniques
- **Validation, Tuning, and Optimization**
  - Assistance in deploying an implementation
- **Support for testing at scale**
  - An unsolved problem, can a CoE add value here?

Create an OFA “Center of Excellence” in Networking

SUMMARY

- **The OFA is in a position to add value**
  - Above and beyond its role in supporting OpenFabrics Software
- **Current interop program is excellent, as far as it goes**
- **By learning from history, we are in a position to offer a much richer program that integrates**
  - Training and support in code development
  - Compliance validation
  - Expert guidance

CALL TO ACTION

- **Engage with the OFA Interoperability Working Group which is beginning to discuss this topic.**

www.openfabrics.org → working groups → IWG working group

- **Subscribe to the mailing list**
- **Check the document archives**
- **Contact the chairperson – Paul Bowden via the link**

OPENFABRICS
ALLIANCE

12th ANNUAL WORKSHOP 2016

THANK YOU