

**OFA XWG Meeting
March 1, 2018
11am Mountain**

1. Roll Call:

Board Members:

At-Large / Bob Noseworthy
Broadcom / Eddie Wai
Cray/Paul Grun
HPE / Andy Riebs
Huawei / Daqi Ren
IBM / Bernard Metzler
Intel / Divya Kolar
Jump Trading / Christoph Lameter
LANL / Susan Coulter
LLNL / Matt Leininger
Mellanox / Gilad Shainer
NetApp / David Dale
Oak Ridge / Scott Atchley
Oracle / David Brean
RedHat/Doug Ledford
Sandia / Mike Aguilar

Also present:

LANL/Jesse Martinez
Mellanox/Bill Lee
OFA/Jim Ryan
Intel/Paul Bowden
LANL/Parks Fields
Intel/Bob Woodruff

2. Approve Board minutes from 3/15 and 3/29

A motion to approved the minutes from 15 March and 29 March were made by: Paul Grun (Cray). The motion was seconded by Susan Coulter (LANL). The minutes were unanimously approved.

3. The new OpenFabrics Alliance Mission Statement is as follows:

The mission of the OpenFabrics Alliance is to accelerate the development and adoption of advanced fabrics for the benefit of the advanced networks ecosystem. The mission is accomplished by: creating opportunities for collaboration among those who develop and deploy such fabrics, incubating and evolving vendor independent open source software for fabrics, and supporting and promoting (the use of) such fabric technologies (software).

A motion to make 2 friendly amendments to the draft Mission Statement were made by Bernard Metzler (IBM) and seconded by Divya Kolar (Intel). The first amendment was

to add in parenthesis (the use of) in the last sentence. A further friendly amendment to add in parenthesis (software) into the last sentence.

Paul Grun (Cray) motioned to approve the new Mission Statement. Divya Kolar (Intel) seconded the motion. The vote was passed unanimously as follows:

Broadcom / Eddie Wai -- yes
Cray/Paul Grun -- yes
HPE / Andy Riebs as a proxy through Susan Coulter (LANL) -- yes
IBM / Bernard Metzler -- yes
Intel / Divya Kolar -- yes
LANL / Susan Coulter -- yes
LLNL / Matt Leininger -- yes
NetApp / David Dale -- yes
Oak Ridge / Scott Atchley -- yes
RedHat/Doug Ledford -- yes
Sandia / Mike Aguilar -- yes

4. SNIA Collaboration with the OpenFabrics Alliance.

A motion was made to approve the collaboration with SNIA was made by Paul Grun (Cray). The motion was seconded by Scott Atchley (ORNL). Susan added that she would like to see the Verbs working group added to the collaboration with SNIA.

The vote passed as follows:

Broadcom / Eddie Wai -- yes
Cray/Paul Grun -- yes
HPE / Andy Riebs as a proxy through Susan Coulter (LANL) -- yes
IBM / Bernard Metzler -- abstained
Intel / Divya Kolar -- yes
LANL / Susan Coulter -- yes
LLNL / Matt Leininger -- yes
NetApp / David Dale -- yes
Oak Ridge / Scott Atchley -- yes
RedHat/Doug Ledford -- yes
Sandia / Mike Aguilar -- yes

A Use/Case description of OpenFabrics Alliance activities and benefits with SNIA is shown below:

SNIA / OFA Work Register
Version 1.0d
Date Initiated: 3/8/2018

This Work Register is created between SNIA and the OFA to formally define the scope, benefits, and deliverables of the alliance. The register helps all organizations coordinate efforts to achieve the stated goals and objectives.

A. A. Alliance Organizations

2

The Storage Networking Industry Association (SNIA), and the OpenFabrics Alliance (OFA).

A. B. Background

SNIA is an organization which develops and promotes standards, technologies, and educational services to empower the management of information in the storage domain. The SNIA Non-Volatile Memory Programming (NVMP) Technical Work Group (TWG) delivers specifications describing the behaviour of a common set of software interfaces that provide access to non-volatile memory (NVM). The TWG goal is to encourage a common ecosystem for NVM-enabled software without limiting the ability to innovate. The OpenFabrics Alliance (OFA) is an open source-based organization that develops, tests, licenses, supports and distributes OpenFabrics Software (OFS). The Alliance's mission is to develop and promote software that enables maximum application efficiency by delivering wire-speed messaging, ultra-low latencies and maximum bandwidth directly to applications with minimal CPU overhead.

A. C. Scope

The specific area of collaboration defined in this Work Register is on remote access to Persistent Memory.

A. D. Alliance Organization Assets

The NVMP TWG is focused on, among other things, describing behavior to be used by applications to access Persistent Memory. The NVMP TWG maintains the following publicly available pieces of collateral relevant to the envisioned collaboration:

- 1. The NVM Programming Model v1.2 SNIA Technical Position*
- 2. The NVM PM Remote Access for High Availability v1.0 white paper*
- 3. Persistent Memory Atomics and Transactions white paper*

The NVM Programming Model Technical Position addresses the ongoing proliferation of new NVM technologies. It defines recommended behaviour of various user space and operating system (OS) kernel components. The specification does not describe any specific API.

The NVM PM Remote Access for High Availability white paper establishes the requirements for the use of RDMA as a transport for remote access to persistent memory in high availability use cases.

The Persistent Memory Atomics and Transactions White Paper describes considerations in developing libraries implementing atomic updates and transactions within the context of byte-addressable persistent memory (PM).

In the context of this proposed collaboration, the API(s) exposed for the purpose of accessing a remote device are out of the NVMP TWG's scope.

OpenFabrics Interfaces *OpenFabrics Interfaces is a project created by the OFA to define interfaces that enable a tight semantic map between user or kernel applications and underlying fabric services. OFI is best described as a framework supporting a collection of libraries and applications used to export fabric services to user or kernel applications. Beside the framework, the key components of OFI are application interfaces, provider libraries, kernel services, daemons, and test applications. The OpenFabrics Interfaces Working Group (OFIWG) is devoted to developing, testing, and distributing an extensible, open source framework that provides access to high-performance fabric interfaces and services. Its primary focus is on user space applications.*

libfabric is a core component of OFI. It is the user space library that defines and exports OFI's user space APIs. The initial focus for libfabric was on distributed and parallel programming applications. Provider libraries, developed by vendors of those libraries, implement the functions described by the API. The libfabric is under control of OFIWG. Each provider library is under the control of its vendor or developer. 3

The Data Storage/Data Access (DS/DA) working group was created to complement OFIWG by focusing on kernel level applications, including kernel space storage applications. The function that corresponds to libfabric is known as kfabric.

A. E. Alliance Benefits

Adoption of remote Persistent Memory as a mainstream technology depends on a) a common understanding among application developers of the behaviors that are required to reliably access remote Persistent Memory, and b) the means for an application to implement those required behaviors by using APIs that are consistent with those required behaviors.

This Work Register is aimed at addressing both points and thus is crucial to driving adoption of remote Persistent Memory.

The role of the NVMP TWG is to describe the behaviors to be exhibited by a consumer of remote Persistent Memory, with the objective of enabling the development of one or more specific APIs.

It is a shared responsibility to define the set of requirements on an API implied by the specified interfaces and/or behaviors.

The role of the OFA is to develop a specific API(s) that conform to the set or requirements developed jointly.

In addition, the OFA is expected to provide feedback to the NVMP TWG on the specifications and behaviors described by the NVMP TWG that result from implementing the set of requirements in a specific API. This collaboration is expected to enable a class of consumers of remote Persistent Memory services that are served by an API meeting their needs while enabling transportability of those applications among multiple types of fabrics.

A. F. Activities

Because remote persistent memory is in the very early stages of definition, the initial objectives for the collaboration are to describe a series of usage models for remote persistent memory, and to define the APIs needed to support those usage models.

Since the HA use case is already well-described, an obvious first step is to review that use case, implement it in the libfabric API, and feed any resulting enhancements back to the NVMP TWG. (Note that there is awareness that work is already underway on this for certain classes of fabrics.)

In parallel, OFI WG and the NVMP TWG will collaborate to enumerate additional use cases. Each relevant use case to be described in a companion whitepaper to the existing HA whitepaper and consisting of a plan or stimulus for concrete API development.

The following activities are envisioned to accomplish this:

- 1. OFI WG and NVMP TWG to institute a series of regular, cross-group engagements at a frequency to be determined by the chairs of the two groups.*
- 2. OFI WG to review the existing HA whitepaper and enhance the existing OFI framework APIs accordingly, subject to availability of volunteer resources a. OFI WG to develop milestones for incorporating the HA use case in the existing OFI framework API(s).*
- 3. OFI WG and the NVMP TWG to collaborate to enumerate further use cases.*
- 4. OFI WG and the NVMP TWG to collaborate to describe each resulting use case, similarly to what NVMP TWG has done for the HA use case. It is expected that these white papers will drive a set of requirements for a resulting API, and that those requirements can be used to drive the development of one or more specific APIs.*
- 5. OFA marketing working group and SSSI to leverage this technical work for cross collaboration in marketing through 2019.*

Limitations 4

OFA and SNIA have differing policies for both how contributions are made, and governance of any outputs of the work. The collaboration must respect those policies. Each organization's contributions are necessarily limited by the availability of volunteer resources.

OFI WG does not control the implementation of the API in any given provider.

Milestones / Dates

OFI WG/SNIA joint table of milestones: Milestone/Deliverables	Timeframe
<i>An update to the requirements described in the existing HA whitepaper based on design of enhancements to the existing libfabric API</i>	<i>Q4'18</i>
<i>A joint enumeration of future use cases to be developed</i>	<i>Q3'18</i>
<i>Support for the collaboration at industry events such as the OFA Workshop, SSSI PM Summit, and other events to be identified.</i>	<i>ongoing</i>
<i>Extensions to the OFI API to support the HA use case</i>	<i>2019</i>