2021 OFA Virtual Workshop

FABRIC SOFTWARE DEVELOPMENT PLATFORM (FSDP)

Tatyana Nikolova tatyana.e.nikolova@intel.com

WHAT IS THE FSDP?

The FSDP is a Hardware Matrix Test Cluster

The FSDP will have hardware from all RDMA IHVs

- InfiniBand – Mellanox with a broad selection of different models/speeds/capabilities (Also in plan custom OEM firmware included as additional variants)
- Omni-Path Architecture – Cornelis
- RoCE – Mellanox, Cavium/QLogic/Marvell, Broadcom, potentially Huawei (subject to changes in current restrictions), Intel
- iWARP – Chelsio, Intel, Cavium/QLogic/Marvell

The FSDP will also include hardware related to RDMA technologies

- NVMe for NVMe over Fabrics testing
- NVDIMM for Remote Persistent Memory over RDMA testing
- GPUs for Peer-to-Peer DMA and GPU direct testing

WHAT IS THE FSDP?

FSDP CI testing will be the third service committed to upstream quality

Intel runs the upstream kernel 0-day testing service

- Builds all kernel patches
- Performs limited boot testing
- Makes no attempt to ensure patches actually work

Google runs Syzkaller testing service

- Runs upstream kernels through syscall validation tests
- Intentionally calls syscalls with known bad data
- Limited support for syscall chains, common in RDMA

The OFA will be running the FSDP CI service

- Runs upstream kernels as well as upstream user space
- Will focus on specific code (RDMA, Peer-2-Peer DMA, etc.)
- Will ensure that code actually runs on the target hardware
- Will utilize an upstream ecosystem to advance tests

BROAD AUDIENCE WITH FLEXIBLE USAGE

Linux Upstream Maintainers	• Automatic, continuous testing of upstream software • Centralized testing and tracking of multiple hardware vendors' products • Development of new software APIs upstream, e.g. GPUDirect
Hardware Vendors*	• On demand testing for IHVs (Mellanox, Intel, Chelsio, Cavium…) • Access to a multi-vendor cluster for development/testing/validation • Logo program, if desired
OS Distros**	• On demand testing for distros (Red Hat, SuSE, OFED, etc.) • Access to a multi-vendor/multi-release cluster for e.g. release testing • Logo program, if desired
ISVs, Applications, Middleware	• On demand testing of specific software • Assist in software development

*served by original OFILP (OpenFabrics Logo Program)
**originally served by the "on-demand" testing program at NMC

WHAT DO YOU GET BY PARTICIPATING IN THE FSDP CI SERVICE?

Upstream kernel community rule:

"If you submit a patch, and it breaks something else, you are responsible for fixing your patch"

The Reality:

- Breakage often caught far too late (months after patch accepted)
- Many hours wasted figuring out which patch caused seemingly unrelated breakage

Proposed Solution:

- Upstream CI catches breakage before patches are officially integrated into upstream code base
- Author will still be working on patch, will be notified of breakage, can easily adapt to fix breakage
- Because fix happens in upstream, trickles down to all distros

Key Takeaways:

- Catch as many bugs introduced by others as possible, and have them fix their patches
- Even when the responsibility to fix the bug falls on your own hands, provides ***months*** more time to fix the bug compared to bugs discovered during distro testing

FSDP DEEP DIVE

FSDP STRUCTURE

FSDP is a cluster managed by a beaker host (beaker-project.org)

- Beaker supports Fedora and Red Hat installs at the moment
- Looking for help to add additional OS support (requires that the OS support automated installs controlled by some sort of control file and a template to create the necessary control files)

Bare metal installs, avoid virtualization effects

Build server with long lived, NFS mountable shares

Direct ssh access to build server and client machines

FSDP STRUCTURE

Git repos for managing the cluster:

- git://github.com/OpenFabrics/fsdp_docs – General cluster documentation
- git://github.com/OpenFabrics/fsdp_setup – Post install setup scripts to configure clients to operate in cluster
- git://github.com/OpenFabrics/fsdp_build – Container definitions for use on build server to allow building for a specific environment
- git://github.com/OpenFabrics/fsdp_tests – Tests available to be run on the FSDP cluster (open for contributions by anyone, but will also be seeded from Red Hat's internal RDMA related tests)

Possibly add containerized infrastructure in the future

USAGE: UPSTREAM CI SERVICE

- **Support the Linux community through a Continuous Integration testing program**
  - Synchronized to, and automatically triggered by, commits to specific git repos
  - A local Continuous Kernel Integration Runner (CKI Runner) daemon patrols for upstream changes
  - Driven by upstream maintainer requested test plans
  - Results reported to an appropriate upstream mailing list

CKI daemon monitors for git commits and kicks off test sequence on changes

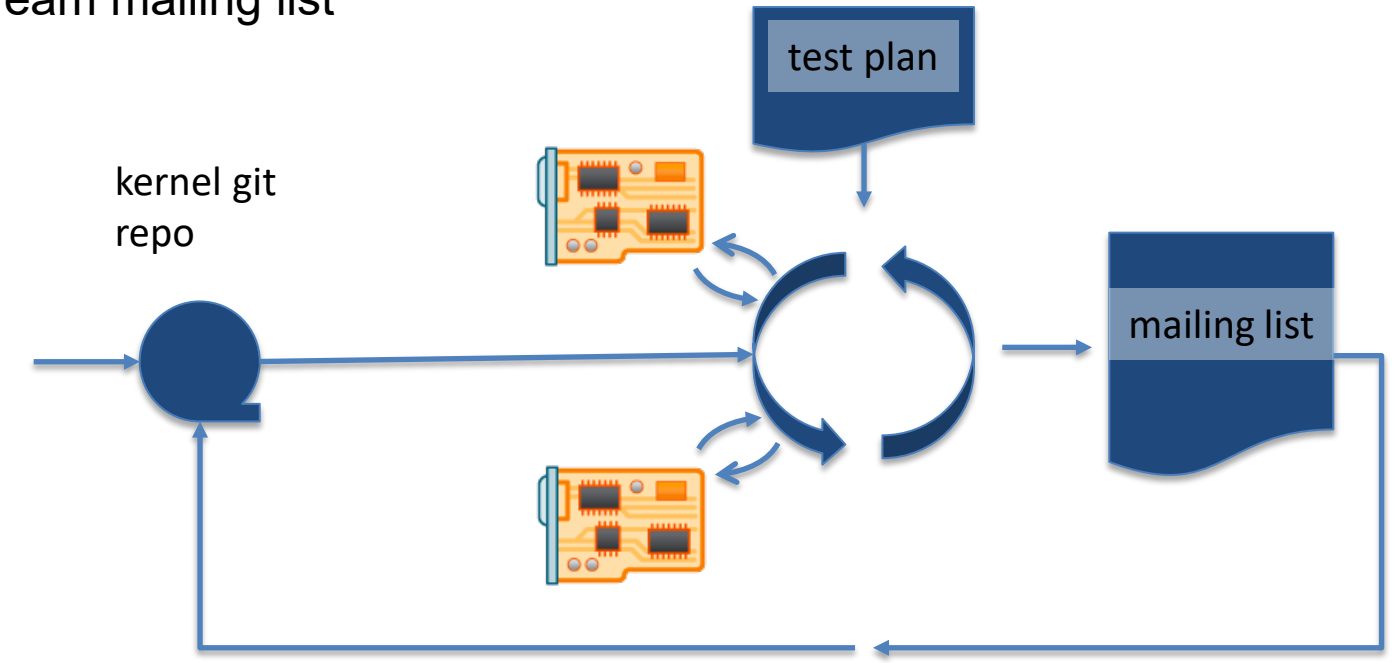

USAGE: ON-DEMAND PROGRAM

- **On-demand program allows for**
  - Development, debug, testing, and design validation
  - May utilize manually initiated automated test runs, or fully manual machine checkouts
  - Checked out machines are an exclusive, dedicated resource for the member with remote ssh access
  - Manually initiated test runs need not be OFA-defined test plans

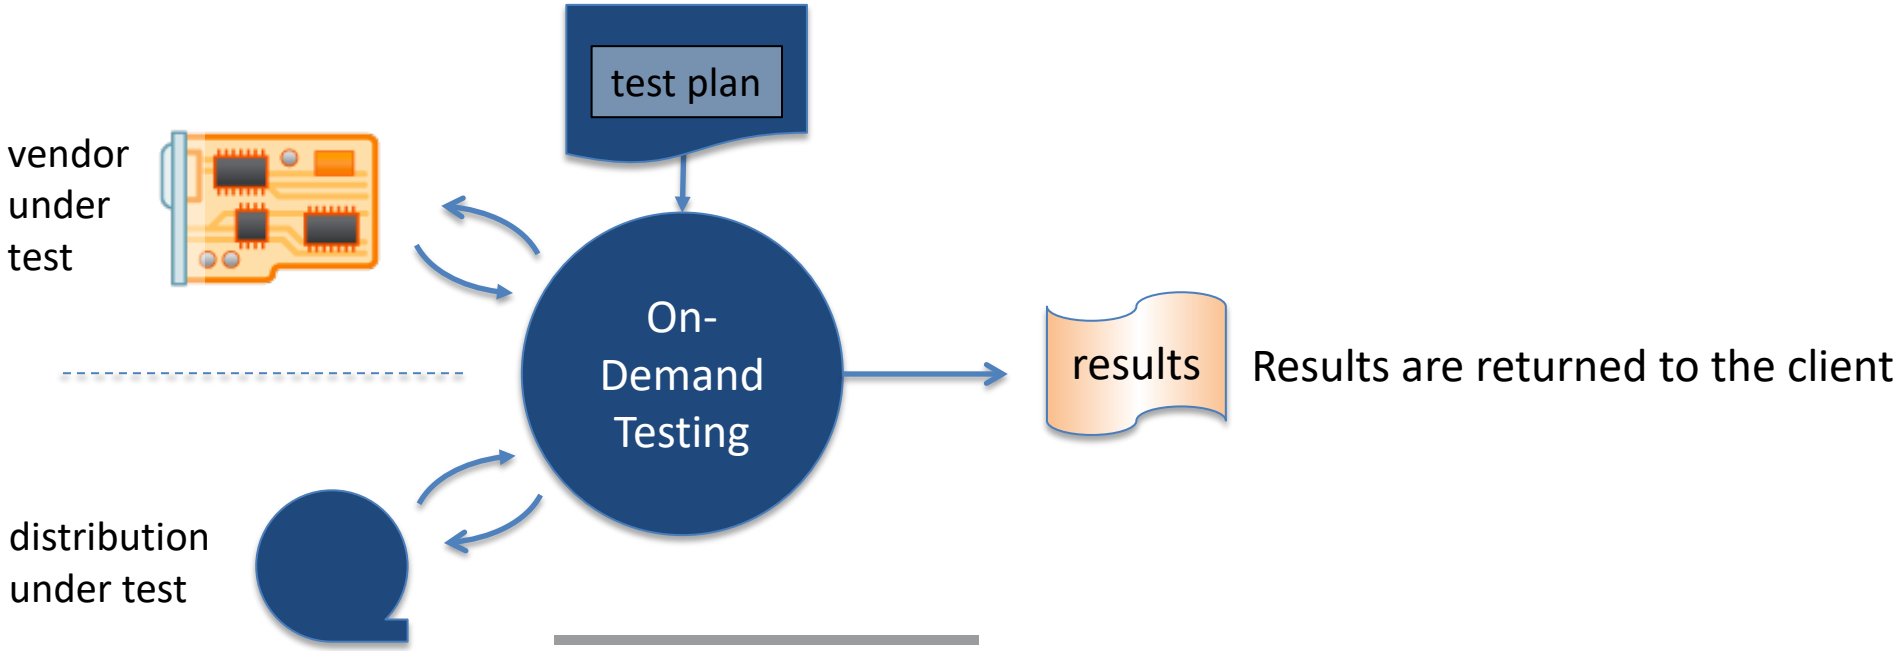

USAGE: LOGO PROGRAM

- **Two possible types of Logos: *Vendor Logo & Distro Logo***
  - Logo tests are run 'on-demand', driven by OFA's test plan as defined by the FSDP Working Group
  - Test plan is executed selectively
  - Run against a defined hardware configuration
  - Run against a specific distribution(s)

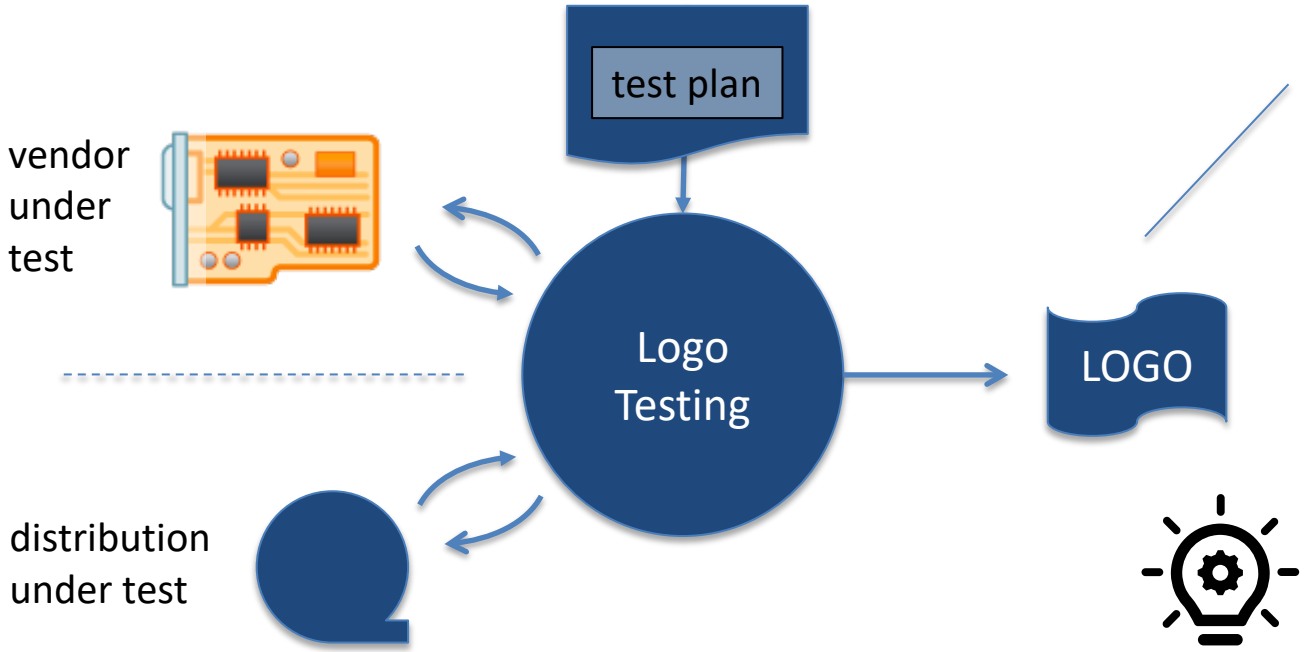

Logo is awarded to Vendor or Distro
Logo Certification includes:
- Test environment
- list of tests executed
- pass/fail results

"Hardware family X is certified to work with RHEL x.x, SLES y.y"
or
"Our distribution supports the following hardware …"

STEPS TO PARTICIPATE IN THE FSDP

PROPOSED MEMBERSHIP LEVELS

Membership Level*	FSDP Participation level
Promoter	• Can be sole chair of FSDP WG • Can appoint a Director to the OFA Board, which then approves appointments to Working Group Chairs/Co-Chairs and Working Group charters
Voting Member	• Can act as Co-Chair for any Working Group and has a vote in Working Groups
Non-Voting Member	• Access to the FSDP cluster and allows the Organization to participate in all Working Groups, however, the Organization will have no vote in Working Groups
Individual	• Free service provided to bona fide upstream developers

- All members are members of the OFA and must abide by the OFA's Intellectual Property Rights Policy
- Have access to the FSDP cluster and must abide by the FSDP Acceptable Use Policy
- Must submit an executed Membership Agreement to membership@openfabrics.org

CALL TO ACTION

- Get an account on the FSDP https://github.com/OpenFabrics/fsdp_docs/blob/main/FSDP_Account_Request.pdf
- Setup your client for access (OpenVPN, ssh are the main requirements)
- Join the FSDP Working Group mailing list (on next slide)

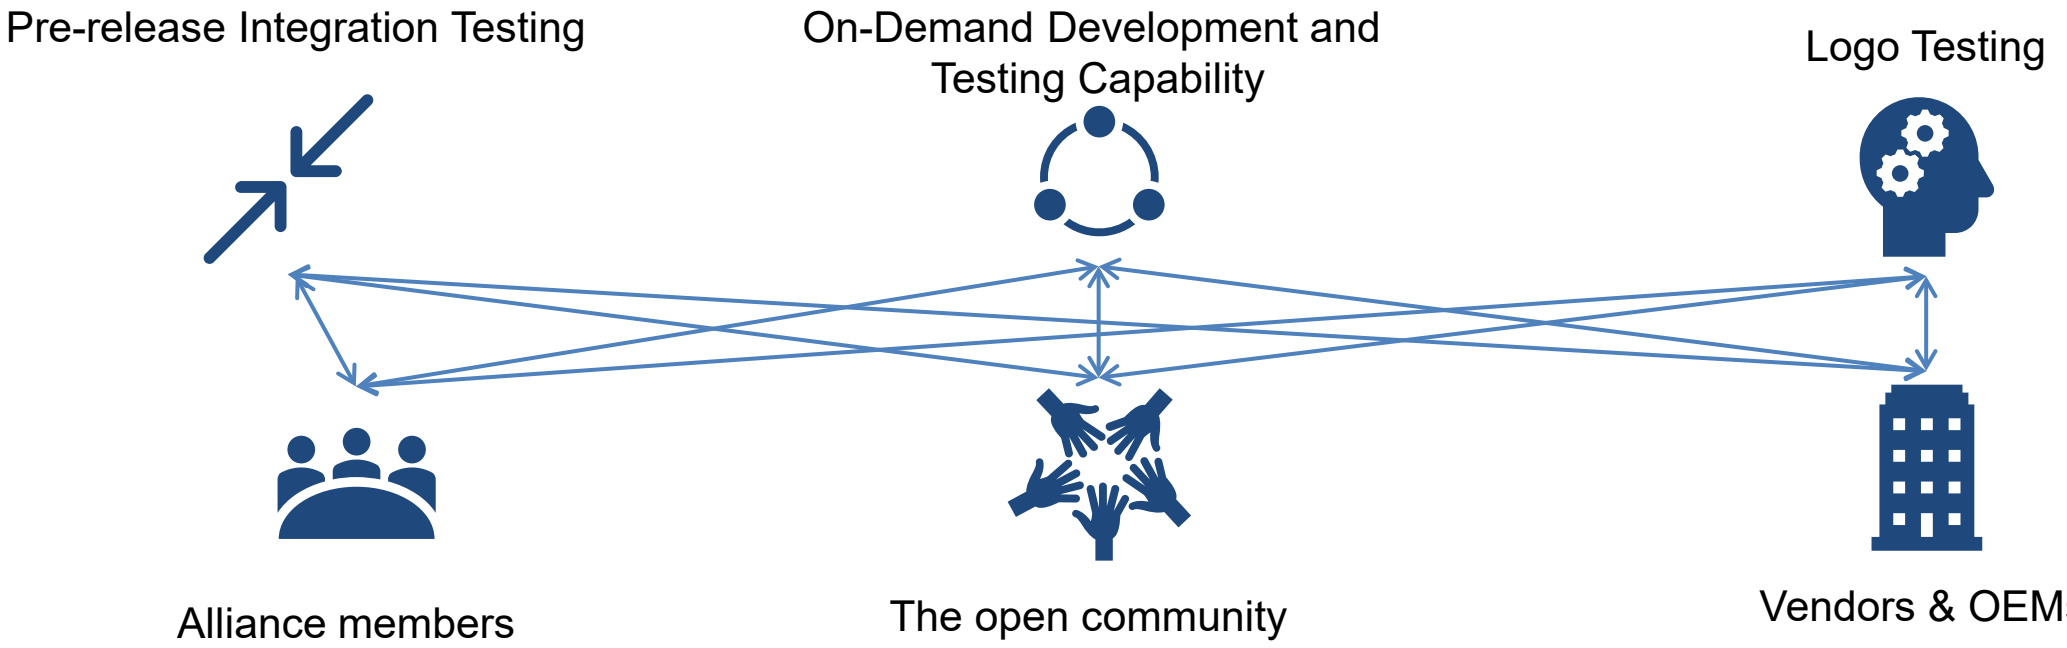

Join OFA and FSDP WG now!

JOIN THE FSDP WORKING GROUP

Oversees the cluster usage and activities

- Arbiter of Acceptable Use Policy violations
- Monitor for members that are wasting resources by checking machines out and then not using them
- Make sure that CI service keeps running smoothly

Logo Program

- Responsible for defining what tests must be passed for any given certification
- Responsible for maintaining the OFA automated test script that IHVs can run as part of a logo attempt
- Will review the results of test runs and approve/deny a logo test

Participation in FSDP WG is open to all, but…

- Chairmanship and voting rights are limited to OFA Voting Members and above
- Send subscribe <email-address> to fsdpwg-requests@lists.openfabrics.org
- fsdpwg@lists.openfabrics.org is the actual mailing list address

ONCE HARDWARE ARRIVES (WHICH HAS ALREADY HAPPENED)

FSDP Working Group Phase 1 – During cluster build

- Get status updates
- Kickstart upstream test repo project
- Early Cluster Access

FSDP Working Group Phase 2 – Once cluster up and running

- Produce webinar series
- Produce FSDP usage tutorial
- Produce FSDP test creation tutorial
- Create Logo program test definitions
- Cluster Generally Available

FSDP Working Group Phase 3 – Maintenance phase

- Routine monitoring and maintenance
- Oversight
- Logo test review/approvals

2021 OFA Virtual Workshop

THANK YOU