



15th ANNUAL WORKSHOP 2019

DRIVING THE INTEROP PROGRAM INTO THE FUTURE

Tatyana Nikolova - Intel, IWG Chair

Paul Grun - Cray, OFA Chair

March 21, 2019

AGENDA

➤ **Why an Interop program at all**

➤ **What the heck is interoperability, and what do we have today?**

- Hardware-based debug + logo program, based at University of New Hampshire Interoperability Lab (UNH-IOL)
- ‘On-demand’ distro testing program, hosted by the New Mexico Consortium (NMC)

➤ **The role of the IWG**

- What the IWG does, how it related to UNH-IOL and NMC

➤ **A Modest Proposal for the next evolution – The OFA Testing Program**

- Build synergy between the existing programs, deliver greater value to participants, the community, and the industry

WHY AN INTEROP PROGRAM

The simple answer - It is a core component of the OFA's mission:

“The mission of the OpenFabrics Alliance (OFA) is to accelerate the development and adoption of advanced fabrics for the benefit of the advanced networks ecosystem.”

The Interop Program is a key element in driving the adoption of advanced fabrics

AN INSTRUCTIVE LOOK IN THE REARVIEW MIRROR

The Basic Objective behind the original Interop Program

- **To build industry confidence in the newly emerging InfiniBand Architecture**

At the time, the OpenFabrics Enterprise Distribution – OFED, was a key industry enabler

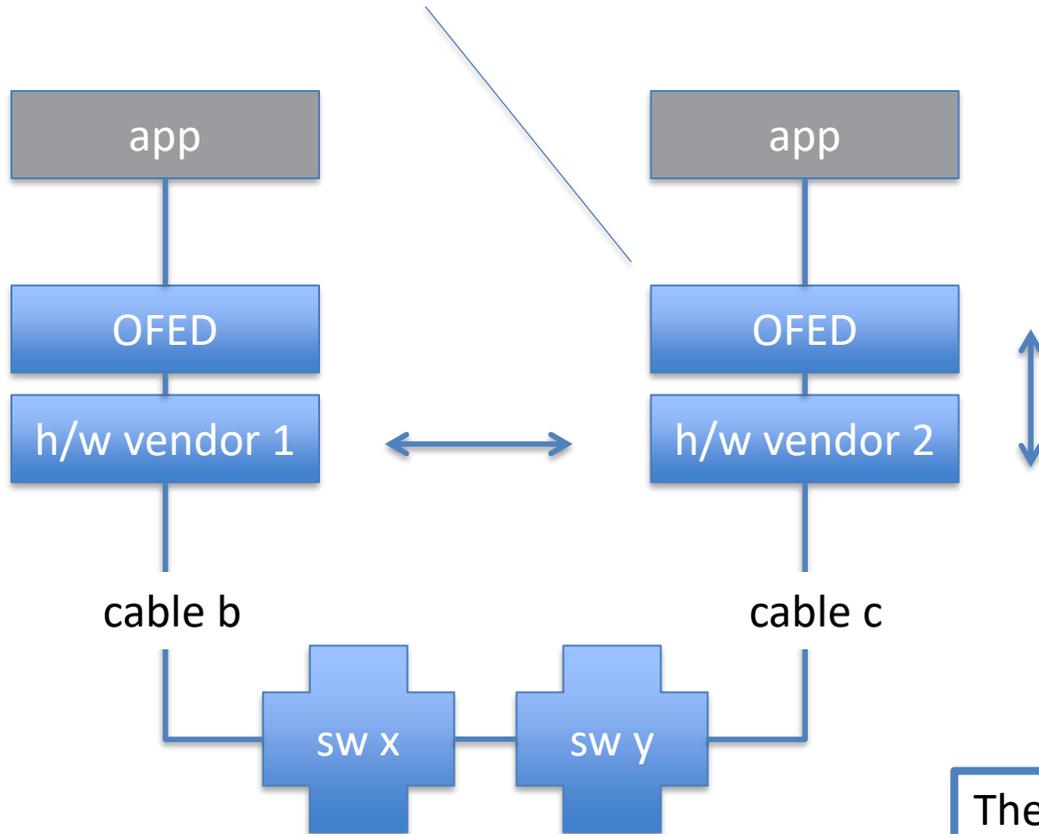
- **It was the only open source software distribution supporting the InfiniBand Architecture**

The purpose of the interop program was to verify that a diverse set of hardware devices:

- **would interoperate among themselves, and**
- **would interoperate with the only existing open source software stack - OFED**

INTEROPERABILITY MEANS:

Historically, the 'distro' was OFED + CentOS



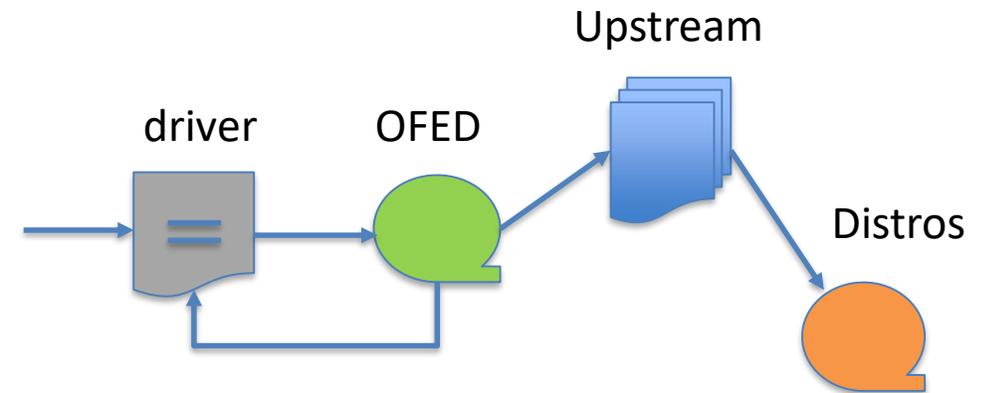
1. Devices interoperate with their peers
 - “Horizontal interoperability”
2. Devices interoperate with the s/w stack
 - “Vertical Interoperability”

The OFA's Interoperability Program (OFILP) has always tested for both

AN ODD THING HAPPENED ALONG THE WAY...

In the beginning,

- new code was pushed to OFED,
- Devices were tested against pre-release OFED,
- Validated drivers were pushed upstream,
- Distros pulled from upstream



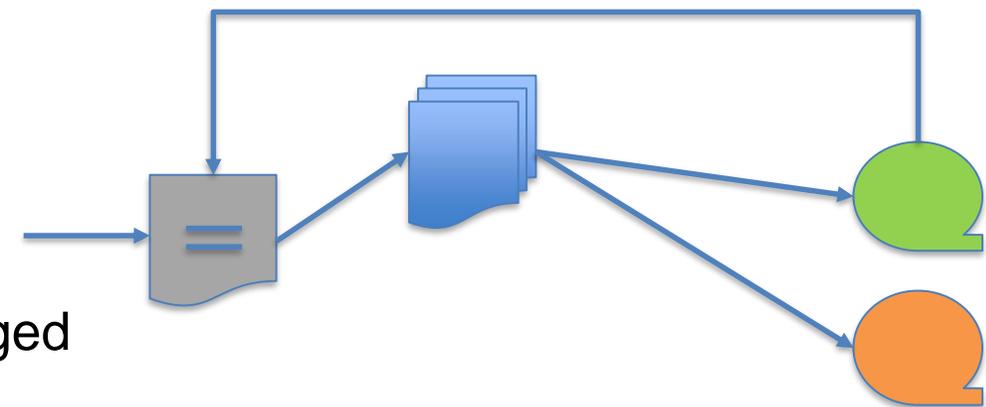
Nowadays, new code is pushed upstream...

- OFED, and the distros, both pull from upstream

This had a profound effect

- Faster release cycles
- The impact of OFED was diminished
- Sacrificed timely interop testing
- Meanwhile, the Interop Program continued unchanged

Clearly, it's time to re-think the Interop Program



THE GOOD & THE BAD (THERE IS NO UGLY)

The Bad

- Testing happens late in the cycle → not aligned with kernel releases
- Diminished value in testing hardware devices against OFED
- No actual testing against commercial distros
- Amortizing the fixed costs across a shrinking field of clients is making the program less and less cost effective to its clients

The Good

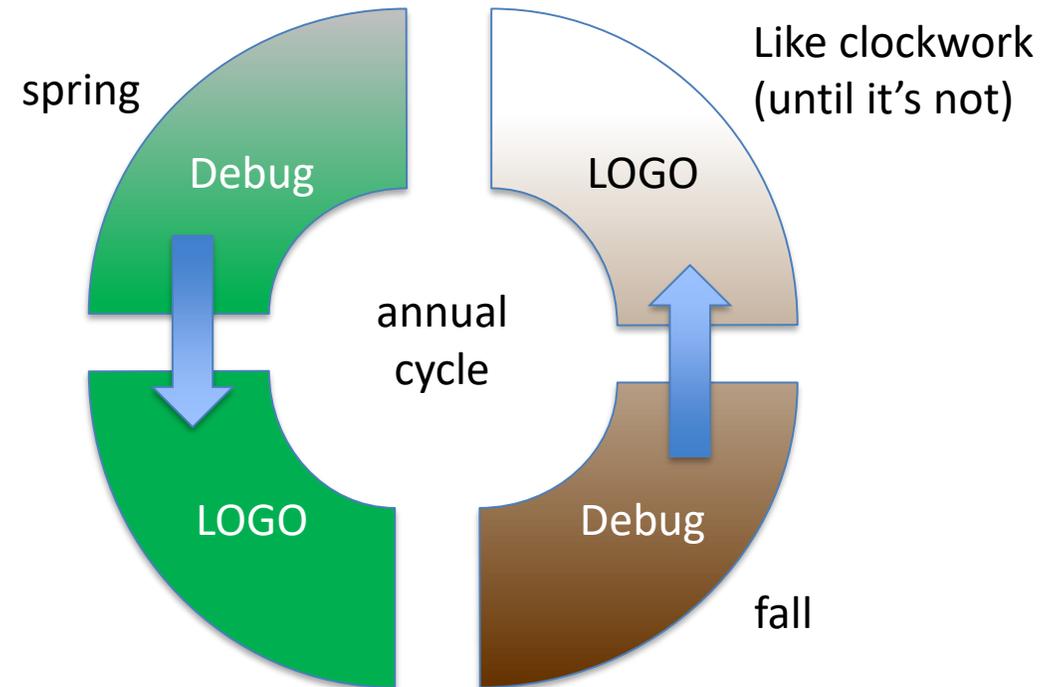
- An opportunity to leverage the new distro testing program → synergy
- An opportunity to add real value to the current kernel release process
- → increased value to program participants, the community, and the industry
- A chance to increase our emphasis on e.g. RoCE, NVMe, libfabric ...

OFILP – TODAY'S PROGRAM

Subscribers pay a fixed base fee, + an incremental fee based on number of devices tested
Program operates on a cost recovery basis, with the OFA serving as the financial backstop

The annual cycle includes two debug events, and two logo events

The single subscription fee bundles two debug events together with two Logo events

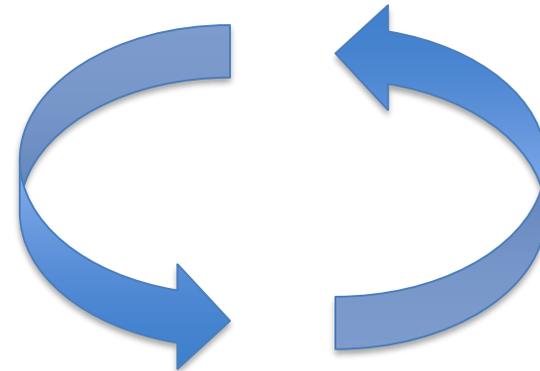


OFILP : OpenFabrics Interoperability Logo Program

DISTRO TESTING PROGRAM

Operates 'on-demand' - pricing is per run

- **Distro submits a release candidate (rc)**
- **The rc is tested according to a defined test plan***
- **Results are returned to the distro**
- **Rinse, lather, repeat**



*currently, it's mostly the same test plan as is used to drive the OFILP

INTEROP GOVERNANCE MODEL - OFA-IWG

Open to participation by all

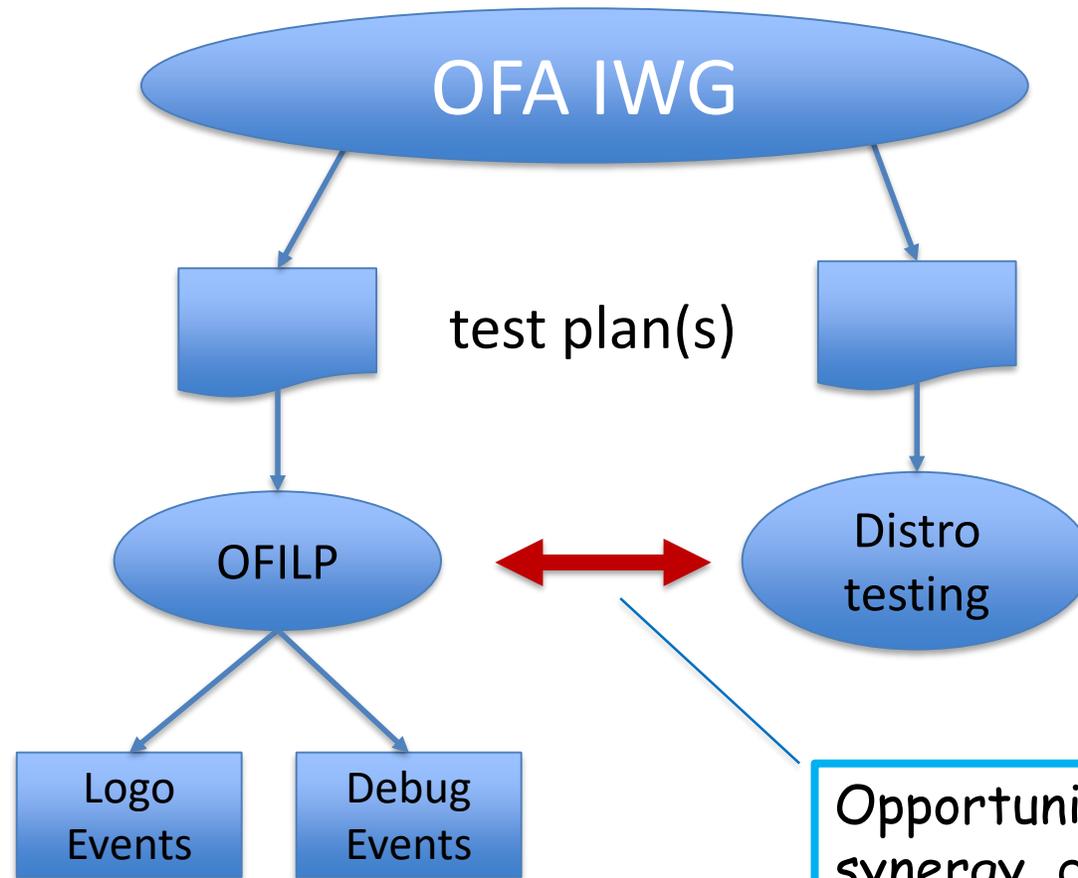
IWG oversees the programs via the Test Plan

OFA Interop Logo Program

- Funded by subscribers
- Confidential

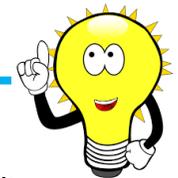
Cadence

- OFILP runs on a regular cadence
- Distro testing is 'on demand'



As of today, the distro testing program is strictly 'debug' (no logo)

Opportunity: Create synergy, deliver greater value



QUESTION



Can we design a program that delivers greater value at lower cost? while serving the needs of:

- Alliance members
- The open community
- Vendors
- OEMs

In short, a program that delivers on the OFA's mission





A MODEST *PROPOSAL*

SCOPE AND STAKEHOLDERS

- **Scope – as always the scope is limited to ‘advanced networks’**
 - Does not include standard networking such as sockets/TCP/IP
 - Does not include compliance to wire specifications
 - Does include testing of software stacks
 - Does allow for testing of interoperability among vendors
- **Direct Stakeholders**
 - Existing subscribers to the OFILP – the Interop program,
 - Kernel maintainers
 - Hardware vendors
 - Distros
 - OpenFabrics Alliance
- **Indirect Stakeholders**
 - OEMs (rely on the IHVs to test hardware)
 - Upstream Linux RDMA community
 - OFA Alliance Members who have a stake in the success of the OFA

SERVING OUR CLIENTS

1.

Support the upstream kernel release cycle

2.

Provide responsive 'On-demand' testing

3.

Continue the Logo Program, but based on standard Linux distributions

STAKEHOLDER WISH LISTS

▪ **Upstream Kernel Pre-release Testing**

- Improve the fidelity of kernel pre-release testing
- Focus on integration testing late in the kernel release cycle (when vendor driver updates have been integrated)
- Ensures that existing drivers interop correctly with a pre-release kernel
- Provides visibility to upstream maintainers into what's been tested by vendors

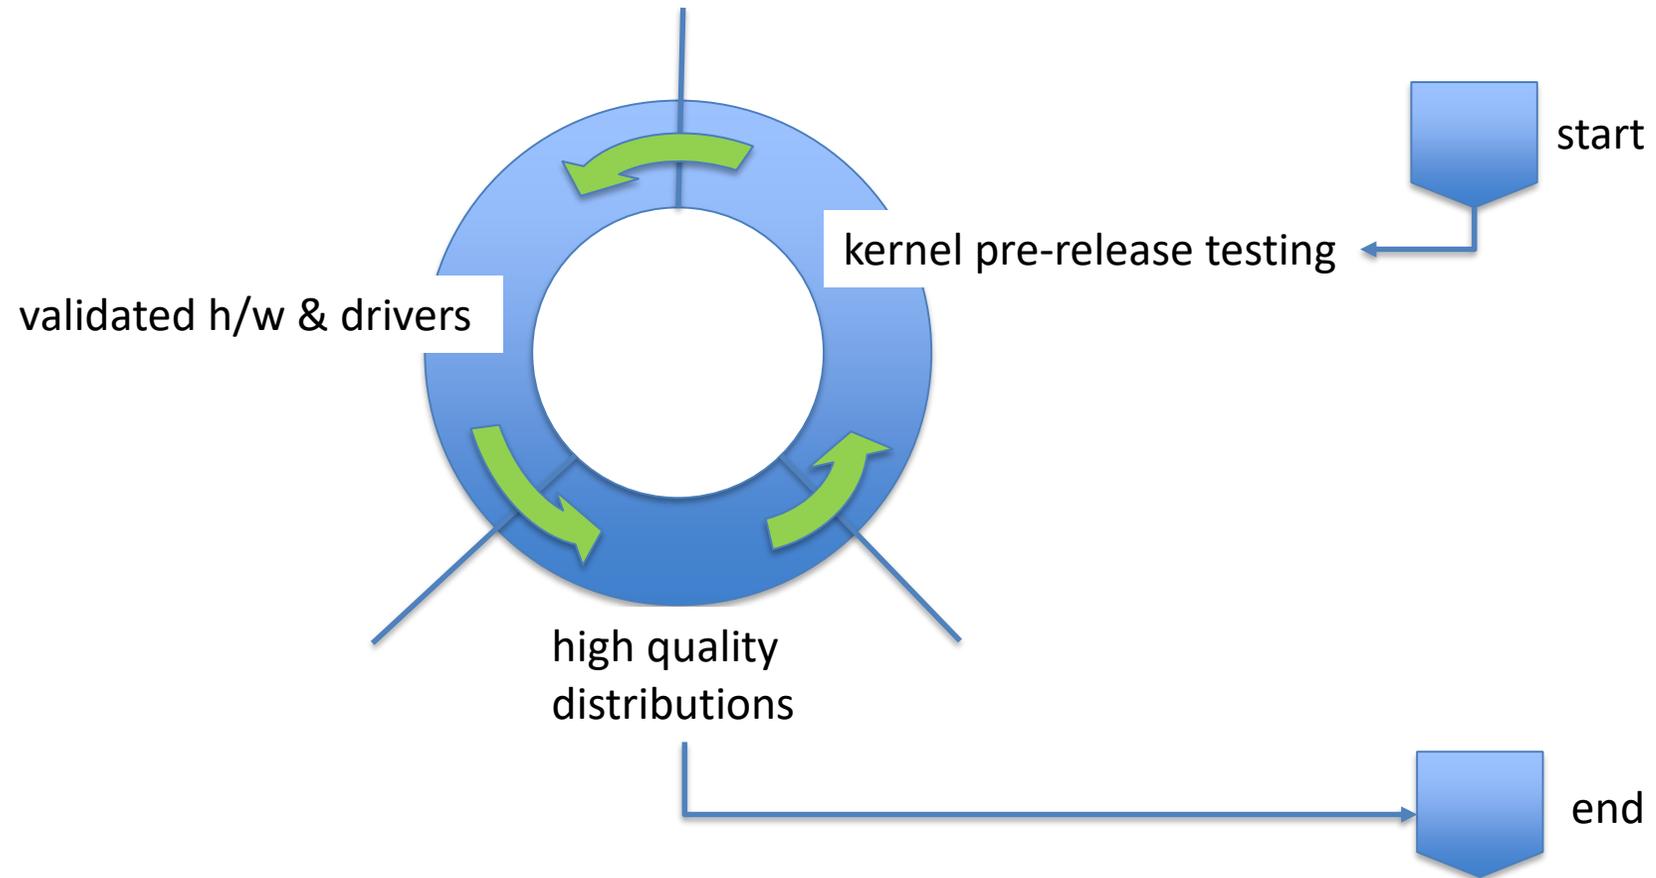
▪ **Hardware Vendors**

- An independent certification program driven by an industry-defined testplan for those who want it
→ a requirement for a logo program
- Reduce testing costs and the ability to test new hardware against other vendors and/or multiple distros
→ a requirement for a 'sandbox' program
- test upstream drivers

▪ **Distros**

- Support for testing release candidates against a stable, current, multi-vendor hardware base
- Test backported drivers – ensure that features that work in the upstream version did not get broken in the Distro
- The opportunity to evaluate vendor backported drivers before inclusion in a GA release
- May value a certification program
→ a requirement for a logo program

A STREAMLINED VIRTUOUS CYCLE



UPSTREAM KERNEL PRE-RELEASE TESTING

■ Why?

- Ensures that existing drivers interop correctly with a pre-release kernel
- Provides visibility to upstream maintainers into what's been tested by vendors
- Creates a fuller picture of device testing during the kernel pre-release process

■ How?

- Provide a known test plan to be used during the kernel pre-release process
- Support integration testing late in the kernel release cycle (when all vendors driver updates have been integrated)
- Provide results to the upstream linux-rdma mailing list for faster response times

TO DO

1. Create a coherent Testing Program

- Build synergy between the existing Interop Program and Distro Testing Program
- Draws from the best features of both
- Deliver greater value to the Linux community, OFA members, hardware vendors, industry, and everybody else

2. Modernize the Interop Program by adding standard distributions

3. Integrate the needs of the Linux community into the Program

4. Respond to demand for 'on-demand' testing

5. Update the current program by modernizing the existing test plan

- The focus is on providing a logo for hardware vendors against a specific set of distro(s)
- OFA continues to own the test plan
- But vendors have some control over which tests are run

6. Add a distro component to the current logo program

- The focus is on providing a validated list of hardware with which a given distro operates correctly
- Against a subset of the testplan benchmarks as defined by the distro

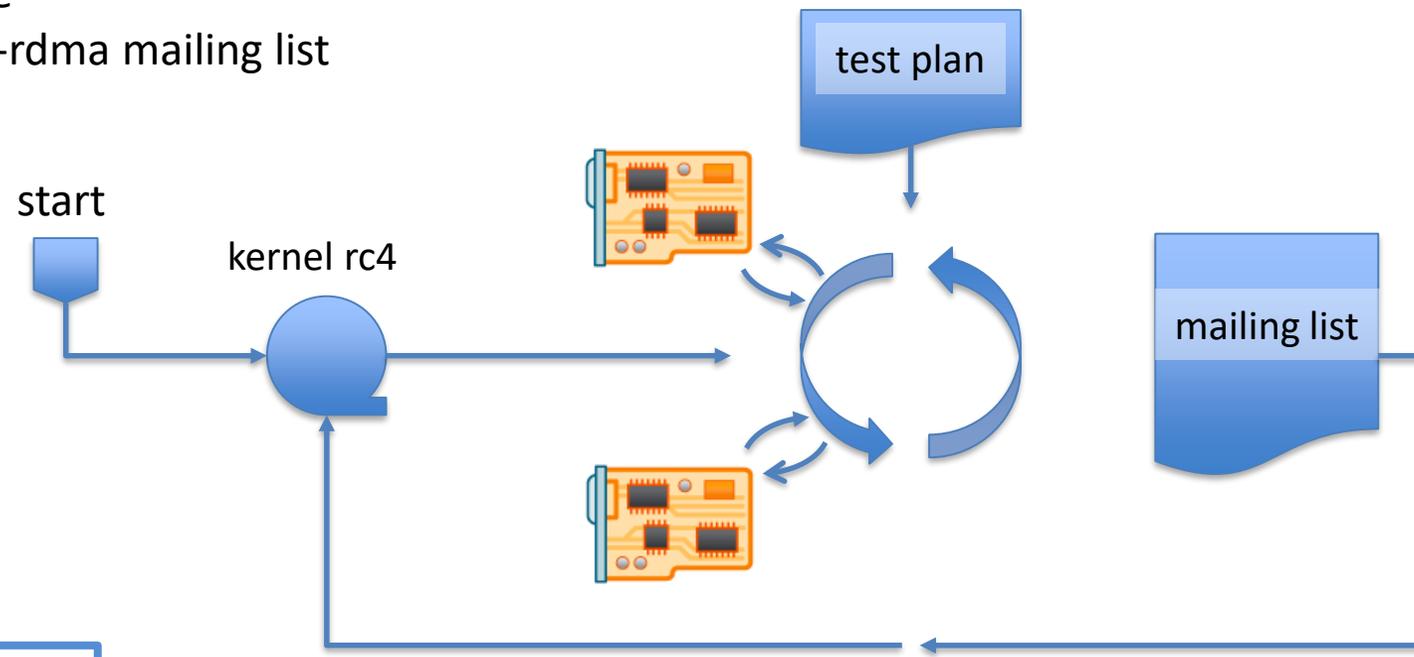
7. Don't preclude compliance testing (e.g. libfabric testing)

8. Test against current, up to date hardware that is constantly maintained

PRE-RELEASE INTEGRATION TESTING

Triggered by kernel release cycle

- Executed late in the integration cycle, rc4 or 5
- Frees the vendors from duplicating tests being run by every other vendor
- Driven by an OFA-defined test plan
- Interactive test and debug cycle
- Coordinated through the Linux-rdma mailing list



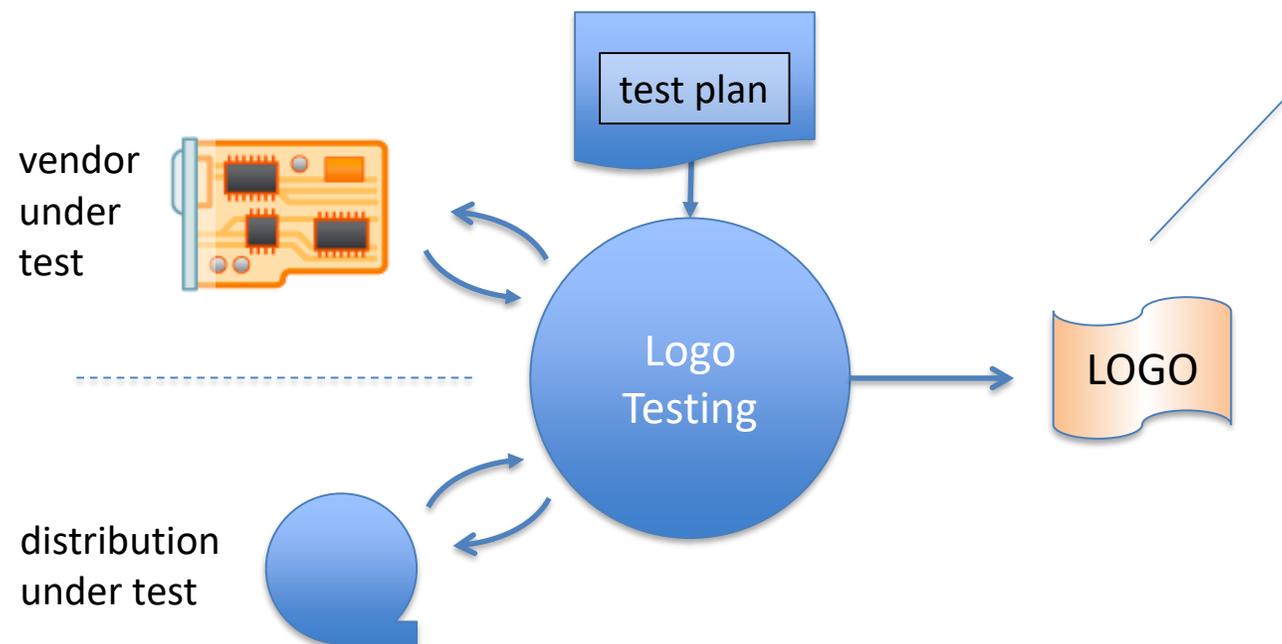
Think of it as collaborative debugging

ON-DEMAND LOGO PROGRAM

Two types of Logos: *Vendor Logo & Distro Logo*

Logo tests are run 'on-demand', driven by OFA's test plan

- Test plan is executed selectively
- Run against a defined ("certified") hardware configuration
- Run against a specific distribution(s)



Logo is awarded to Vendor or Distro
Certification includes:

- Test environment
- list of tests executed
- pass/fail results

"Our new hardware is certified to work with RHEL x.x, SLES y.y"
or
"Our distribution is supported by the following hardware"

PRICING – TBD!

- **Currently, program costs are bundled, and include participation in both the debug event and the subsequent logo event**
 - An all or nothing proposition
- **One possibility: eliminate the current subscription model**
- **Base-level participation as a function of membership level**
- **Logo pricing is priced per run**

NEXT STEPS

→ Meet our new IWG Chair – Tatyana Nikolova ←

- Come to the BoF to discuss the strawman proposal – Pros and Cons
- Make the proposal better!
- Work with our testing vendors on the program details
 - New Mexico Consortium - NMC
 - University of New Hampshire Interoperability Lab – UNH-IOL
- Detailed proposal to the OFA's Board of Directors in the next few months

Join the Interoperability Working Group
Be an active part in driving this forward



15th ANNUAL WORKSHOP 2019

THANK YOU

Tatyana Nikolova

Paul Grun