OFA-IWG Interoperability Test Plan: Suggested Modifications for RoCE

The following notes are based on v1.46 of the OFA-IWG Interoperability Test Plan. This is by no means a complete and fully specified list of changes, but rather a set of discussion points (along with some questions) which I hope to use to drive appropriate modifications & enhancements to the RoCE test portions.

Page 19: RDMA stress

I think that we need to simultaneously stress RoCE/IB traffic and IP level Ethernet traffic. This will ensure that the Ethernet and IB/RDMA portions of the hardware, drivers & libraries work properly together. Perhaps we can use something such as `uperf` (`www.uperf.org`) to generate standard Enet traffic in conjunction with various RDMA traffic generators; perhaps `dapltest?` or just various tests from the `perftest` and `librdmacm-utils` rpms?

Page 24, Table 21

For clarification, how is IPoCE different from just standard IP over the RCA? Or is this referring specifically to IP traffic over a CEE fabric?

Page 26, 1.2 Homogeneous vs. Heterogeneous

We need to add and allow mixed system architectures for x86_64 and ppc64 interoperability. This also implies mixed endianness between those systems.

Section 9.2, Operating System

For the IBM Power system, RHEL6.x will have to be used. There are no CentOS, Scientific Linux or Ubuntu distributions for Power platforms.

Section 13.1.3, TI testing (iSER)

I don't foresee any iSER or RDS testing over RoCE. And maybe this point is already moot given the discussion of reduced ULP testing.

General Enhancements

- RDMA_CM
  We would like to see explicit rdma_cm tests, particularly for processor-heterogeneous (x86_64/ppc64) setups.
- For heterogeneous testing, it would be good to use tests that are network byte order aware and the results of data transfers are subsequently validated.
- Test rsockets over RoCE
- Test IPv6
- Test bonding over RoCE Enet interfaces

Future Considerations

- RoCE & VLANs. In the Mellanox implementation, a VLAN ID is stored as the 12th and 13th bytes of the GID. Is that compatible across vendors?
- Validating Priority Flow Control (802.1Qbb)