



2023 OFA Virtual Workshop

USING THE FSDP FOR CI ON RDMA HARDWARE

Doug Ledford

Chair of the OFA, Former upstream RDMA maintainer, Manager in Red Hat's kernel department



OPENFABRICS
ALLIANCE

WHAT THE FSDP IS

A SMALL TEST MATRIX CLUSTER

■ What does test matrix mean?

- As a small cluster, it's not suitable for scale out testing
- The machines in the cluster are not configured like typical machines in a HPC cluster
 - Multiple interfaces, using different fabrics, e.g. a single machine with iWARP, RoCE, InfiniBand, and OmniPath Architecture all present (node-03 and node-04)
 - A table of what machines have what hardware present can be found in the docs repo:
https://github.com/OpenFabrics/fsdp_docs/blob/main/test_nodes.md
- Very few identical pieces of hardware. Instead, the cluster emphasizes different hardware. Although hardware is installed in pairs so that you can always test homogeneous first and then test heterogeneous later.
- General documentation: https://github.com/OpenFabrics/fsdp_docs/blob/main
- Accounts:
 - Any employee or contractor for an OFA member company can request an account
 - Any open source contributor can also request an account

MANUAL TESTING

How to use the FSDP

1. Provision a machine or two
2. Select the OS you want installed
3. Beaker will install the OS to a standard configuration
4. Your default root password or ssh public key as configured in your account will be the default on the machines you provision
5. Use ssh to log into the machines and perform your manual testing

```
[dledford@Fedora-WSL ~]$ ssh node-03
```

```
Warning: Permanently added 'node-03' (ED25519) to the list of known hosts.
```

```
Activate the web console with: systemctl enable --now cockpit.socket
```

```
Register this system with Red Hat Insights: insights-client --register
```

```
Create an account or view all your systems at https://red.ht/insights-dashboard
```

```
Last login: Tue Apr 11 22:29:41 2023
```

```
/usr/bin/which: no rhts-reboot in (/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin)
```

```
[root@node-03 ~]$ en
```

```
lom_1: Link UP, Interface UP 10.12.2.15/24
```

```
mlx5_roce: Link UP, Interface UP 172.31.40.15/24 <- Subnet for TCP/IP traffic
```

```
mlx5_roce.43: Link UP, Interface UP 172.31.43.15/24 <- Subnet for RoCE traffic
```

```
mlx5_roce.45: Link UP, Interface UP 172.31.45.15/24 <- Subnet for RoCE traffic
```

```
cxgb4_iw: Link UP, Interface UP 172.31.50.15/24 <-\
```

```
cxgb4_iw.51: Link UP, Interface UP 172.31.51.15/24 <--- All subnets usable for iWARP
```

```
cxgb4_iw.52: Link UP, Interface UP 172.31.52.15/24 <-/
```

```
[root@node-03 ~]$ ib
```

```
mlx5_ib0: Link UP, Interface UP 172.31.0.15/24 <-\
```

```
mlx5_ib0.8002: Link UP, Interface UP 172.31.2.15/24 <--\
```

```
mlx5_ib0.8004: Link UP, Interface UP 172.31.4.15/24 <--- Various MTUs on different nets, otherwise all the same
```

```
mlx5_ib0.8006: Link UP, Interface UP 172.31.6.15/24 <-/
```

```
[root@node-03 ~]$ opa
```

```
hfi1_opa0: Link UP, Interface UP 172.31.20.15/24 <-\
```

```
hfi1_opa0.8022: Link UP, Interface UP 172.31.22.15/24 <--- All subnets usable
```

```
hfi1_opa0.8024: Link UP, Interface UP 172.31.24.15/24 <-/
```

```
[root@node-03 ~]$
```

USER INITIATED AUTOMATED TESTING

How to use the FSDP

1. Check out the fsdp_tests repo: https://github.com/OpenFabrics/fsdp_tests
2. Either use existing tests or add new tests of your own design
3. Make sure you've installed the beaker-client package (this might require downloading directly from the upstream project site: <https://beaker-project.org/download.html>)
4. Configure your `~/beaker_client/config` file
5. Use the `bkr` command line program to initiate test jobs

```
cat ~/.beaker_client/config

# Base URL of the Beaker server (without trailing slash!)
HUB_URL = "https://beaker.ofa.iol.unh.edu/bkr"

# Hub authentication method
# Some installs use krbv instead
AUTH_METHOD = "password"

# Username and Password
USERNAME = "dledford"
PASSWORD = "YourPasswordHere"

# If the OFA used kerberos for the beaker server, these would be useful
# Kerberos principal. If commented, default principal obtained by kinit is used.
#KRB_PRINCIPAL = "host/$HOSTNAME"

# Kerberos keytab file.
#KRB_KEYTAB = "/etc/krb5.keytab"

# Kerberos service prefix. Example: host, HTTP
#KRB_SERVICE = "HTTP"

# Kerberos realm. If commented, last two parts of domain name are used.
# Example: MYDOMAIN.COM.
#KRB_REALM = "EXAMPLE.COM"

# Kerberos credential cache file.
#KRB_CCACHE = ""

# SSL CA certificate to verify the Beaker server's SSL certificate.
# By default, uses the system-wide CA certificate bundle.
# The OFA cert can be downloaded from beaker.ofa.iol.unh.edu/ca.crt
#CA_CERT = "/etc/beaker/OFA_CA.pem"
#CA_CERT = "/etc/pki/tls/cert.pem"
```

```
[dledford@Fedora-WSL demo (main *)]$ bkr workflow-xslt --job-xml
./multihost.xml -d "Fedora Rawhide" -f FedoraRawhide -a x86_64 -v Server -s
node-03.ofa.iol.unh.edu -c node-04.ofa.iol.unh.edu --whiteboard="FSDP
Multihost Test" -e iw -i cxgb4
```

Generating Beaker XML

```
Job config:      ./multihost.xml
XSLT template:  multihost.xsl
Job name:       FSDP demo of a simple RDMA test
Whiteboard:    FSDP Multihost Test
Job arguments:
- family: FedoraRawhide
- variant: Server
- distro: Fedora Rawhide
- arch: x86_64
- server: node-03.ofa.iol.unh.edu
- client: node-04.ofa.iol.unh.edu
- kickstart:
- env_driver: cxgb4
- env_network: iw
```

```
** Sending job to Beaker ...
```

```
Success
```

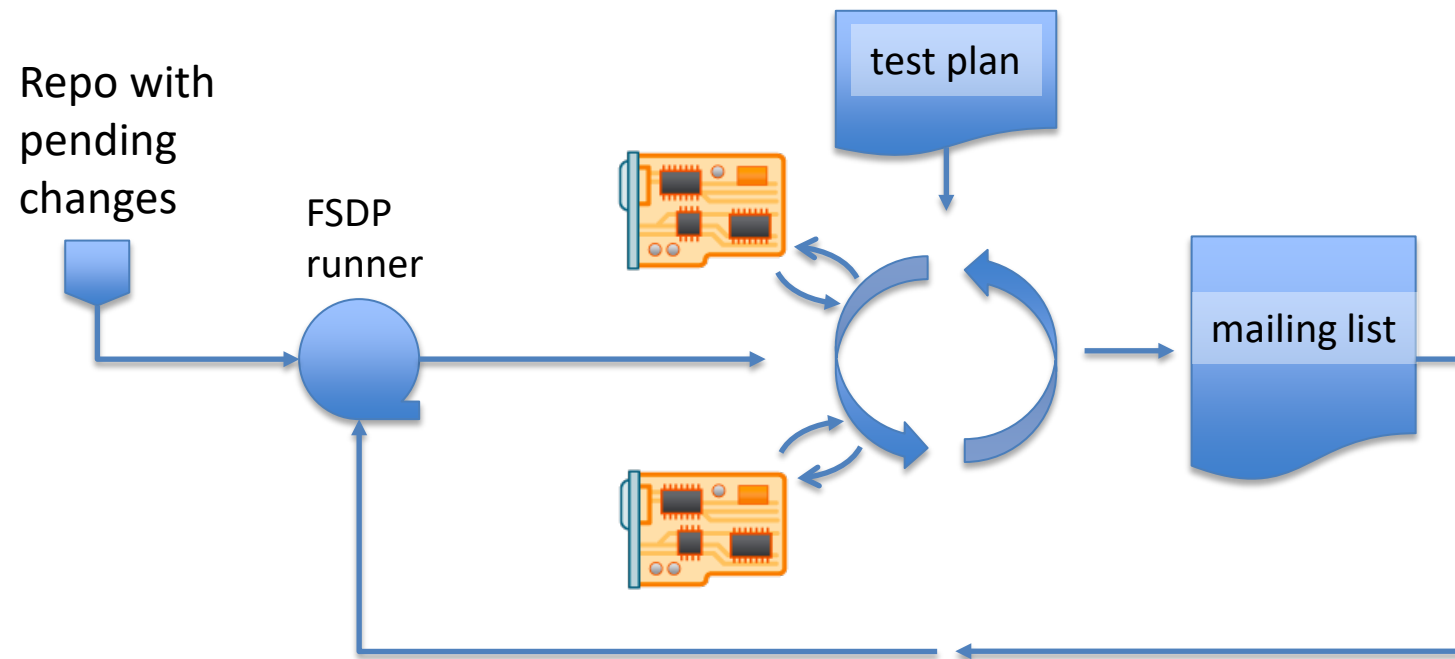
```
** Submitted: ['J:1071']
```

```
[dledford@Fedora-WSL demo (main *)]$
```

CI INITIATED AUTOMATED TESTING

How to use the FSDP

1. The Red Hat CKI infrastructure is configured to watch your upstream repo (see <https://gitlab.com/cki-project/pipeline-data/-/blob/main/ofa.yaml>)
2. For a kernel repo, only the git_url is needed
3. For a user space repo, you need to specify the kernel you want the user space tests run on in the git_url, and then the user space repo in watch_url and watch_branch
4. The test to run is specified in the test_set element, and can be modified by adding new test sets to the kpet-db repo (see <https://gitlab.com/redhat/centos-stream/tests/kernel/kpet-db>)
5. The test set element in kpet-db then points to the FSDP test repo



THE FSDP IS AVAILABLE FOR USE

▪ Steps to get your project added for CI

- Follow the instructions in https://github.com/OpenFabrics/fsdp_docs to create an account in the FSDP
- Update ofa.yaml in <https://gitlab.com/cki-project/pipeline-data> and file a merge request
- Update index.yaml and any other needed files in <https://gitlab.com/redhat/centos-stream/tests/kernel/kpet-db> and file a merge request
- Create and test any new tests in https://github.com/OpenFabrics/fsdp_tests, using manual initiated automated testing, and once the tests work as expected, file a merge request
- Contact FSDP Working Group mailing list fsdpwg@lists.openfabrics.org and request review for your merge requests

▪ Special considerations for user space CI tasks

- The CKI infrastructure is really built for kernel CI (CKI stands for Continuous Kernel Integration)
- As a result, the CKI pipelines only build the kernel (and always build a kernel, which is why the git_url of a kernel tree is required for user space pipelines)
- The first task of a user space test suite in the fsdp_tests repo needs to be a synchronized task to build the needed user space software and install it, then after synchronizing on the install, the subsequent testing can begin



OPENFABRICS
ALLIANCE

2023 OFA Virtual Workshop

THANK YOU

Doug Ledford