



LNET Selftest

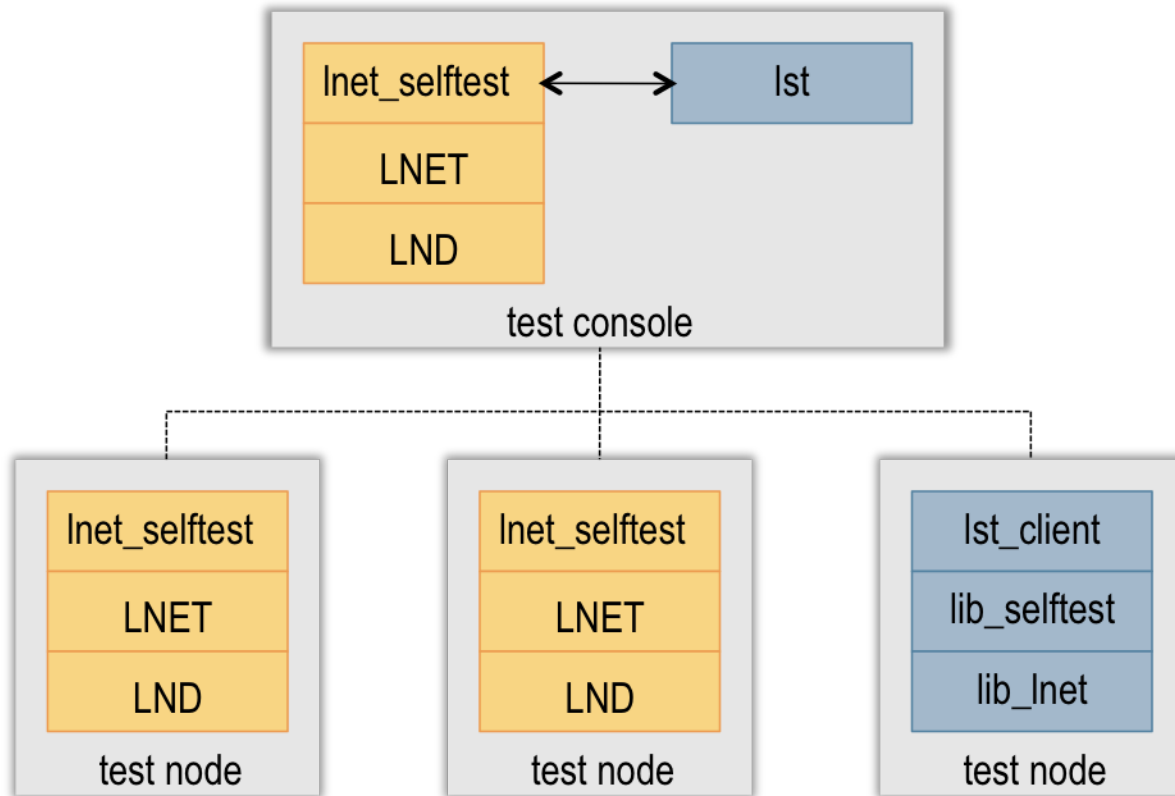
- Isaac Huang
- Lustre Group



Introduction and Overview

- Easy to setup and use
 - > Setup: only configure LNET properly
 - > Use: everything done at test console
- Ping, bulk data read/write (with data integrity checks)
- Extensible
- Selftest in userspace
- Use cases
 - > Smoke test
 - > Performance measurement

Components



Key Concepts

- Session
- Console
- Group
- Batch
- Test

A Sample script

```
# setup test session: add nodes, add tests
lst new_session rw
lst add_group servers 192.168.10.[8,10,12-16]@tcp
lst add_group readers 192.168.1.[1-253/2]@o2ib
lst add_group writers 192.168.1.[2-254/2]@o2ib
lst add_batch bulk_rw
lst add_test --batch bulk_rw --from readers --to servers brw read check=simple size=1M
lst add_test --batch bulk_rw --from writers --to servers brw write check=full size=4K
# run test and gather stats
lst run bulk_rw
lst stat servers
lst end_session
```

A Typical Smoke test

```
lst new_session soak
```

```
lst add_group s 192.168.0.[1-51]@tcp
```

```
lst add_group c 192.168.0.[100-150]@tcp
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from c --to s brw write check=full size=4K
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from c --to s brw write check=full size=64K
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from c --to s brw write check=full size=512K
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from c --to s brw write check=full size=1M
```

```
.....
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from c --to s ping
```

```
lst add_test --loop 8192 --concurrency 8 --distribute 1:1 --from s --to c ping
```

```
lst run
```

```
lst stat c s
```

Sample 'show_error' output

```
lsh show_error clients
```

```
clients
```

```
12345-192.168.1.15@tcp:
```

```
[Session: 1 brw errors, 0 ping errors]
```

```
[RPC: 20 errors, 0 dropped,
```

```
12345-192.168.1.16@tcp:
```

```
[Session: 0 brw errors, 0 ping errors]
```

```
[RPC: 1 errors, 0 dropped, Total 2 error nodes in clients
```

Concepts Continued

- Distribution defines mapping between test client group and test server group
- Loop equals iteration
- Concurrency specifies the number of test instances running simultaneously

A Typical Bandwidth Test

```
lst new_session perf
```

```
lst add_group s 10.0.0.1@o2ib
```

```
lst add_group c 10.0.0.2@o2ib
```

```
lst add_batch bulkr
```

```
lst add_test --batch bulkr --loop 8192 --concurrency 8  
--distribute 1:1 --from c --to s brw read check=none  
size=1M
```

```
lst run bulkr
```

```
lst stat c
```

Sample 'stat' output

[LNet Rates of c]

[W] Avg: 667 RPC/s Min: 667 RPC/s Max: 667 RPC/s

[R] Avg: 1341 RPC/s Min: 1341 RPC/s Max: 1341 RPC/s

[LNet Bandwidth of c]

[W] Avg: 0.11 MB/s Min: 0.11 MB/s Max: 0.11 MB/s

[R] Avg: 673.71 MB/s Min: 673.71 MB/s Max: 673.71 MB/s

[LNet Rates of c]

[W] Avg: 677 RPC/s Min: 677 RPC/s Max: 677 RPC/s

[R] Avg: 1365 RPC/s Min: 1365 RPC/s Max: 1365 RPC/s

[LNet Bandwidth of c]

[W] Avg: 0.11 MB/s Min: 0.11 MB/s Max: 0.11 MB/s

[R] Avg: 687.96 MB/s Min: 687.96 MB/s Max: 687.96 MB/s

Selftest in Userspace

- Test console must run in kernel
- Create session at console:
 - > `lst new_session userland_test`
- Userspace test nodes join session:
 - > `lstclient --sesid CONSOLE_NID group NAME`
- Set up tests and batches at console
- Run test and gather stats

Q & A



• he.huang@sun.com

