LNET Router Resiliency and Tuning

John Fragalla
HPC Principal Engineer
Seagate Technology

SEAGATE

Lustre User Group Conference, April 13-15, 2015

Agenda

- Configuration Goal
- Benchmark Setup
- LNET Router Configuration/Testing
  - Fine Tuning and testing 4MB I/O
- Summary
- References

Configuration Goal

- Extend Data Access to Ethernet connected clients from an InfiniBand based Lustre System
  - 1 IB Lustre File System, IB Clients for performance, Ethernet Routed Clients for additional data access
- Resiliency and Performance of IB Clients were higher priority than optimizing performance of Routed Clients

Benchmark Setup

- Lustre Storage System based on ClusterStor 9000 – IB System
  - Lustre MGS NID: 172.18.54.79@o2ib0
  - Lustre MDS NID: 172.18.54.80@o2ib0
  - 4 OSS Nodes with 4 GridRaid OSTs, 2 HA Pairs
- InfiniBand Clients (16 Clients)
- Ethernet Routed Clients (4 Clients), 10GigE Interface
  - 172.19.62.[170,172,174,178]@tcp0
- LNET Routers (4 Configured)
  - InfiniBand ib0 / Ethernet eth20 Interface (10GigE)
  - 172.18.62.162@o2ib / 172.19.62.162@tcp
  - 172.18.62.164@o2ib / 172.19.62.164@tcp
  - 172.18.62.166@o2ib / 172.19.62.166@tcp
  - 172.18.62.168@o2ib / 172.19.62.168@tcp

Lustre Software and Kernel Versions

- Server Side, tested two Lustre Versions
  - Seagate Lustre `2.1.0.x5-148` on Kernel `2.6.32-220.7.1.el6.lustre.4226.x86_64`
  - Seagate Lustre `2.5.1.x6-107` on Kernel `2.6.32-431.17.1.x2.0.43.x86_64`
- Clients (LNET Router, Ethernet, and InfiniBand Clients)
  - Kernel: `2.6.32-431.23.3.el6.x86_64`
  - Client version:
    - 1.8.9, 2.1, 2.4.3 (with Server version 2.1)
    - 2.4.3 and 2.5.1 (with Server version 2.5.1)

Client Mount Options

- checksum=0
- max_rpc_in_flight = 256
- max_dirty_mb=256
- max_pages_per_rpc=256 (1MB I/O)
- Mount Command

```
mount –t lustre 172.18.54.79@o2ib0:172.18.54.80@os2ib0:/lustre /mnt/lustre
```

Initial LNET Router Configuration and Results

Initial LNET Router Configuration – Server

Using Reference Materials

- /etc/modprobe.d/lnet.conf

```
options lnet dead_router_check_interval=60
options lnet routes="/etc/lustre/routes.dat”
options lnet networks=“o2ib0(ib0)”
options lnet avoid_asym_router_failure=1
options ko2iblnd timeout=100 peer_timeout=130
options ko2iblnd credits=2048 ntx=2048
options ko2iblnd peer_credits=126 concurrent_sends=63
peer_buffer_credits=128
options lnet check_routers_before_use=1
options lnet live_router_check_interval=60
options lnet router_ping_timeout=50
options lnet large_router_buffers=1025 small_router_buffers=16384
```

- /etc/lustre/routes.conf

```
tcp0 172.18.62.[162,164,166,168]@o2ib0
```

Initial LNET Router Configuration – LNET Router

Using Reference Material

- /etc/modprobe.d/lnet.conf

```
options lnet dead_router_check_interval=60
options lnet avoid_asym_router_failure=1

options ko2iblnd timeout=100 peer_timeout=130
options ko2iblnd credits=2048 ntx=2048
options ko2iblnd peer_credits=126 concurrent_sends=63
peer_buffer_credits=128

options lnet check_routers_before_use=1
options lnet live_router_check_interval=60
options lnet router_ping_timeout=50
options lnet large_router_buffers=1025 small_router_buffers=16384
options lnet ip2nets="tcp0(eth20) 172.19.62.*;o2ib0(ib0)
172.18.62.*;o2ib0(ib0) 172.18.62.[162,164,166,168]"
options lnet routes="o2ib0 172.19.62.[162,164,166,164]@tcp0"
```

Initial LNET Router Configuration – Routed Ethernet Clients

Using Reference Material

- /etc/modprobe.d/lnet.conf

```
options lnet dead_router_check_interval=60
options lnet networks=”tcp(eth20)”
options lnet avoid_asym_router_failure=1

options lnet check_routers_before_use=1
options lnet live_router_check_interval=60
options lnet router_ping_timeout=50
options lnet large_router_buffers=1025 small_router_buffers=16384
options lnet routes="o2ib0 172.19.62.[162,164,166,164]@tcp0”
```

- /etc/modprobe.d/ko2iblnd.conf

```
options ko2iblnd timeout=100 peer_timeout=130
options ko2iblnd credits=2048 ntx=2048
options ko2iblnd peer_credits=126 concurrent_sends=63
peer_buffer_credits=128
```

Initial LNET Router Configuration – IB Clients

Using Reference Material

- /etc/modprobe.d/lnet.conf

```
options lnet networks="o2ib0(ib0)"
```

- /etc/modprobe.d/ko2iblnd.conf

```
options ko2iblnd timeout=100 peer_timeout=130
options ko2iblnd credits=2048 ntx=2048
options ko2iblnd peer_credits=126 concurrent_sends=63
peer_buffer_credits=128
```

Initial Results

- Using

```
options ko2iblnd timeout=100 peer_timeout=130
options ko2iblnd credits=2048 ntx=2048
options ko2iblnd peer_credits=126 concurrent_sends=63
peer_buffer_credits=128
```

  - Configuration is a bit more complex using these options
  - On Client, saw a performance drop including these options in lnet.conf compared using a separate ".conf file" for the parameters.
- Routed and Non-Routed Clients: Inconsistent lnet.conf parameters
  - Simulating a router failure, and I/O is only occurring on Ethernet Routed Clients
  - Writes Complete, but did not start the Read Phase, IOR was in a hung state
- Outcome
  - Remove the ko2iblnd parameters and configure all clients to have consistent lnet.conf parameters, regardless if the Client is IB or Routed Ethernet

Consistent lnet.conf Parameters Results - Resiliency Testing

- Observation on Writes during Router Failure
  - When 1 Router Fails, leaving 3 Up, I/O pauses from 2 to 4 minutes, then resumes.
  - Reasons for the range of 2 to 4 minutes pause is the timing from the LNET parameters:

```
options lnet dead_router_check_interval=60
options lnet live_router_check_interval=60
options lnet router_ping_timeout=50
```

- Observation on Reads during Router Recovery
  - When the failed Router from Write Phase is recovered during Read Phase, Read never pauses, and throughput increases within 1 minute after the Router is recovered
- Observation of Reads during Router Failure
  - Failing 1 out of 4 routers during Read phase, I/O paused for approximately 60 seconds, than I/O resumes
- Tested multiple Failures with I/O on IB Clients and Routed Ethernet Clients,
  - MDT Failover, OST Failover, Multiple Router Failure/Recovery
  - No impact on I/O and job completes

Phase 2 LNET Router Configuration: Weighted versus Round Robin

LNET Router Configuration –Routed Clients

- /etc/modprobe.d/lnet.conf

```
options lnet dead_router_check_interval=60
options lnet networks=”tcp(eth20)”
options lnet avoid_asym_router_failure=1

options lnet check_routers_before_use=1
options lnet live_router_check_interval=60
options lnet router_ping_timeout=50
options lnet large_router_buffers=1025 small_router_buffers=16384
```

Client 1 -> options lnet routes="o2ib0 1 172.19.62.162@tcp0;o2ib0 2 172.19.62.164@tcp0;o2ib0 3 172.19.62.166@tcp0;o2ib0 4 172.19.62.168@tcp0”

Client 2 -> options lnet routes="o2ib0 1 172.19.62.164@tcp0;o2ib0 2 172.19.62.166@tcp0;o2ib0 3 172.19.62.168@tcp0;o2ib0 4 172.19.62.162@tcp0"

Client 3 -> options lnet routes="o2ib0 1 172.19.62.166@tcp0;o2ib0 2 172.19.62.168@tcp0;o2ib0 3 172.19.62.162@tcp0;o2ib0 4 172.19.62.164@tcp0"

Client 4 -> options lnet routes="o2ib0 1 172.19.62.168@tcp0;o2ib0 2 172.19.62.162@tcp0;o2ib0 3 172.19.62.164@tcp0;o2ib0 4 172.19.62.166@tcp0"

Results using Weighted Routing

Peak Performance

```
[root@sjsc-73 ~]# pdsh -w sjsc-[73-76] cat /proc/sys/lnet/peers | sort
sjsc-73: 172.19.62.162@tcp         315    up    -1    32    32    32  -242  -474 160510904
sjsc-73: 172.19.62.164@tcp          33    up    -1    32    32    32    32    31 0
sjsc-73: 172.19.62.166@tcp          53    up    -1    32    32    32    32    31 0
sjsc-73: 172.19.62.168@tcp          42    up    -1    32    32    32    32    31 0
sjsc-73: nid                      refs state  last   max   rtr   min    tx   min queue
sjsc-74: 172.19.62.162@tcp          37    up    -1    32    32    32    32    31 0
sjsc-74: 172.19.62.164@tcp         324    up    -1    32    32    32  -247  -421 160513704
sjsc-74: 172.19.62.166@tcp          44    up    -1    32    32    32    32    31 0
sjsc-74: 172.19.62.168@tcp          43    up    -1    32    32    32    32    31 0
sjsc-74: nid                      refs state  last   max   rtr   min    tx   min queue
sjsc-75: 172.19.62.162@tcp          37    up    -1    32    32    32    32  -341 0
sjsc-75: 172.19.62.164@tcp          53    up    -1    32    32    32    32    31 0
sjsc-75: 172.19.62.166@tcp         311    up    -1    32    32    32  -242  -433 161558992
sjsc-75: 172.19.62.168@tcp          44    up    -1    32    32    32    32    31 0
sjsc-75: nid                      refs state  last   max   rtr   min    tx   min queue
sjsc-76: 172.19.62.162@tcp          47    up    -1    32    32    32    32  -213 0
sjsc-76: 172.19.62.164@tcp          39    up    -1    32    32    32    32    31 0
sjsc-76: 172.19.62.166@tcp          37    up    -1    32    32    32    32    31 0
sjsc-76: 172.19.62.168@tcp         296    up    -1    32    32    32  -226  -467 152117240
sjsc-76: nid                      refs state  last   max   rtr   min    tx   min queue

Max Write: 1081.74 MiB/sec (1134.29 MB/sec)
Max Read:  3056.69 MiB/sec (3205.17 MB/sec)
```

Results using Round Robin Routing

Peak Performance

```
[root@sjsc-73 ~]# pdsh -w sjsc-[73-76] cat /proc/sys/lnet/peers | sort
sjsc-73: 172.19.62.162@tcp          45    up    -1    32    32    32    14   -24 18875664
sjsc-73: 172.19.62.164@tcp          39    up    -1    32    32    32    13   -26 19924312
sjsc-73: 172.19.62.166@tcp          44    up    -1    32    32    32    13   -18 19924312
sjsc-73: 172.19.62.168@tcp          38    up    -1    32    32    32    13   -37 19924312
sjsc-73: nid                       refs state  last   max   rtr   min    tx   min queue
sjsc-74: 172.19.62.162@tcp          68    up    -1    32    32    32     5   -16 28313496
sjsc-74: 172.19.62.164@tcp          48    up    -1    32    32    32     6   -14 27264848
sjsc-74: 172.19.62.166@tcp          53    up    -1    32    32    32     6   -23 27264848
sjsc-74: 172.19.62.168@tcp          57    up    -1    32    32    32     6  -134 27264848
sjsc-74: nid                       refs state  last   max   rtr   min    tx   min queue
sjsc-75: 172.19.62.162@tcp          58    up    -1    32    32    32     6   -16 27264848
sjsc-75: 172.19.62.164@tcp          53    up    -1    32    32    32     6   -15 27264848
sjsc-75: 172.19.62.166@tcp          56    up    -1    32    32    32     6   -15 27264848
sjsc-75: 172.19.62.168@tcp          57    up    -1    32    32    32     6  -134 27264848
sjsc-75: nid                       refs state  last   max   rtr   min    tx   min queue
sjsc-76: 172.19.62.162@tcp          50    up    -1    32    32    32     9   -24 24118904
sjsc-76: 172.19.62.164@tcp          52    up    -1    32    32    32     9   -36 24118904
sjsc-76: 172.19.62.166@tcp          43    up    -1    32    32    32    10   -28 23070256
sjsc-76: 172.19.62.168@tcp          55    up    -1    32    32    32     8   -80 25167552
sjsc-76: nid                       refs state  last   max   rtr   min    tx   min queue

Max Write: 3434.50 MiB/sec (3601.33 MB/sec)
Max Read:  3048.00 MiB/sec (3196.06 MB/sec)
```

Comparison in Resiliency Testing

Weighted versus Round Robin

- When simulating Router Failure using weighted routes during IOR Write phase, I/O halts on average of only 60 seconds, and resumes on the healthy routers.
- 2-4 minutes is the average time to resume Write I/O on the client using the failed router as primary before resuming I/O on the backup Router
- Weighted Routing is more deterministic in failure and less impact on all I/O being effected during Router Failure, compared to Round Robin
- Round Robin has a 3x higher Write performance over Weighted Routing, on average
- Weighted Routes can be useful grouping set a of clients to a subset of Routers, doing a combination of Weighted and Round Robin

Phase 3 LNET Router Configuration Using ksocklnd

ksocklnd LNET Routing Configuration

- ksocklnd module is used to fine tune TCP connections
- /etc/modprobe.d/ksocklnd.conf Clients and Router Side

```
options ksocklnd peer_credits=16
options ksocklnd tx_buffer_size=0
options ksocklnd rx_buffer_size=65536
```

- /etc/modprobe.d/ksocklnd.conf Server Side

```
options ksocklnd peer_credits=32
```

Using ksocklnd

Benefits

- Minor performance improvement on writes and reads for Routed Clients

```
Max Write: 3547.06 MiB/sec (3719.36 MB/sec)
Max Read:  3104.70 MiB/sec (3255.51 MB/sec)
```

- No impact on performance for IB Clients
- Better Router Failover time during Resiliency testing, approximately 60 to 120 seconds to resume I/O

Phase 4 LNET Router Configuration and Results testing 4MB I/O

Client Mount Options

- checksum=0
- max_rpc_in_flight = 256
- max_dirty_mb=256
- max_pages_per_rpc=1024 (4MB I/O)
- Mount Command

```
mount –t lustre 172.18.54.79@o2ib0:172.18.54.80@os2ib0:/lustre /mnt/lustre
```

LNET Routing Results testing 4MB I/Os

■Using 4MB I/O had interesting results during simulated Router Failures in this specific system setup
■No Router Failure, performance had a positive impact for IB clients and I/O completes for Ethernet Routed and IB Clients

■Using 4MB I/O on all clients, Simulating Router Failure during Write Phase during IOR, I/O is hung, never resumed
■IB Clients set to use 4MB I/O, but Ethernet Clients set to use 1MB I/O, Simulating router failure during Write Phase, I/O hung and never resumed (same state)
–MTU size, enabling or disabling small/large buffer credits, enabling or disabling ksocklnd had no improvement on 4MB I/O resiliency simulating router failure

–Changing back max_pages_per_rpc to 256 for all clients (1MB I/O) demonstrated no issue for I/O under failed Router test cases (write or read)

New Development: Solution Identified using 4MB I/O with Routers under Failed Condition

- A Server side Patch developed fixes the 4MB I/O hung threads (multi Bulk I/O in general) under Failed Router conditions
- This patch is newly developed and will be pushed back to the community soon (in process)
- Test Case that demonstrates the patch resolved the issue
  - Started IOR and failed routed 1 minute into the test
  - Partial Writes resume after 1 minute, full writes resume after 3 minutes
  - Before Read Phase, recovered Router
  - During Read Phase, failed Router, Read pause for 1 minute, then resumes
  - Recovered Router and full Read bandwidth increased within 1 minute

Summary

- When configured correctly, using LNET routers to extend data access from an IB Lustre system to Ethernet clients works reliably
- Consistent LNET configuration is required across all clients
- ksocklnd settings improves I/O during Router failure for Ethernet Clients
- Multi-Bulk I/O patch addresses the hung I/O issue during Router Failures
- Round Robin gives better performance during normal operation, but Weighted Routing provides significant benefits during router failures.

References

- Seagate Acknowledgement
  - Lustre Development Team
  - Bill Loewe, Rob Roy, Scott Milk, Karl Merritts, Nathan Rutman, Chris Walker, Torben Kling Petersen, Carlos Thomaz, Rex Tanakit
- Online Materials
  - https://cug.org/proceedings/attendee_program_cug2012/includes/files/pap166-file2.pdf
  - http://cdn.opensfs.org/wp-content/uploads/2011/11/LUG-2012-Day-2-930.pdf
  - http://docs.cray.com/cgi-bin/craydoc.cgi?mode=View;id=S-2483-22;idx=books_search;this_sort=release_date%20desc;q=lnet;type=books;title=OpenFabrics%20Interconnect%20Drivers%20for%20Cray%20XT%20Systems
  - https://build.hpdd.intel.com/job/lustre-manual/lastSuccessfulBuild/artifact/lustre_manual.xhtml#dbdoclet.50438272_73839

Thank You

john.fragalla@seagate.com