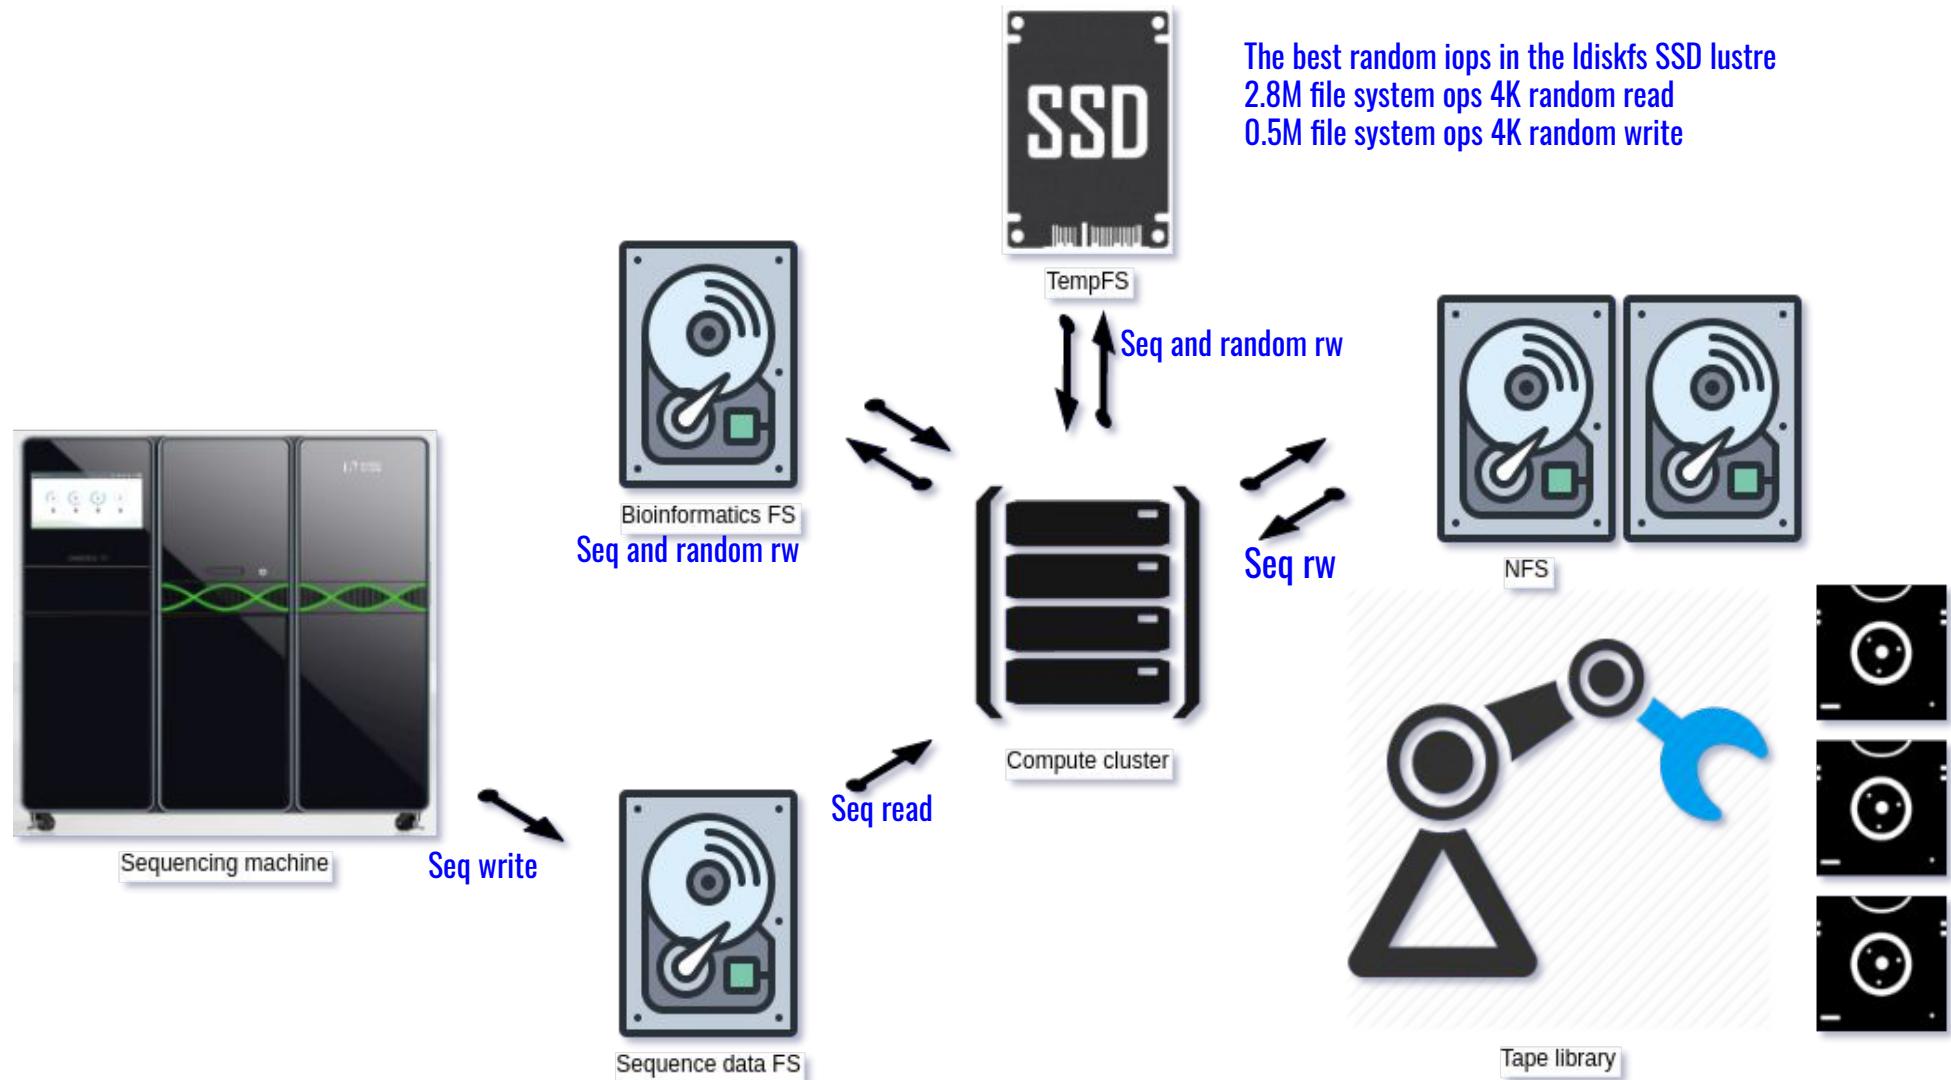


# 系统边界拾遗

2021.12 李焱 Homer Li

# 目录

1. 工作流
2. 以太网问题
3. 其它Tips



# PHY signal

Oversize\_pkts\_phy, err\_lane\_x\_phy

pcs\_symbol\_err\_phy, crc\_errors, crc\_errors\_phy

Tx\_discards\_phy, rx\_symbol\_err\_phy, Jabbers\_phy, rx\_errors

Optical Transceiver compatibility/Cable/PCIE issue

# 监控资源和设置引起的问题

TCPBacklogDrop | ListenDrops

ListenDrops | ListenOverflows

SyncookiesSent | SyncookiesRecv

TimeWaitOverflow

TCPToZeroWindowAdv

AbortOnMemory | AbortFailed

PruneCalled | MemoryPressures

PFMemallocDrop

rx\_fifo\_errors/rx\_discards\_phy/rx\_missed\_errors

port.rx\_dropped/rx\_discards



# 网络抖动

过于复杂的链接类似一个混沌系统

TCPAbortOnData | AbortOnSyn

TCPPartialUndo | TCPFullUndo

TCPDSACKOldSent | TCPDSACKUndo

TCPSynRetrans

LostRetransmit | RetransFail

SACKOfRecv | SACKOfSent | SACKDiscard

# Lustre client 内存问题1

TcpExtPFMemallocDrop	128362	0.0
TcpExtPFMemallocDrop	201640	0.0
TcpExtPFMemallocDrop	1107595	0.0
TcpExtPFMemallocDrop	71210	0.0

Call Trace:

```

<IRQ> [] dump_stack+0x19/0x1b
[<ffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
[<ffffffff88e6ed01>] ? sk_filter_trim_cap+0x31/0xe0
[<ffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffff04d8036>] ? iptable_filter_hook+0x36/0x80 [iptable_filter]
[<ffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450
[<ffffffff88a18e18>] __alloc_pages_current+0x98/0x110
[<ffffffff889e5748>] kmalloc_order+0x18/0x40
[<ffffffff88a243d6>] kmalloc_order_trace+0x26/0xa0
[<ffffffff88a28361>] kmalloc+0x211/0x230
[<fffffff0557dca>] bnx2x_frag_alloc.isra.58+0x2a/0x40 [bnx2x]
[<fffffff0558de7>] bnx2x_rx_int+0x227/0x1950 [bnx2x]
[<fffffff055c47d>] bnx2x_poll+0x1dd/0x260 [bnx2x]
[<fffffff88e545bf>] net_rx_action+0x26f/0x390
[<fffffff888a5695>] __do_softirq+0xf5/0x280
[<fffffff88f9642c>] call_softirq+0x1c/0x30
[<fffffff8882f715>] do_softirq+0x65/0xa0
[<fffffff888a5a15>] irq_exit+0x105/0x110
[<fffffff88f97876>] do_IRQ+0x56/0xf0
[<fffffff88f8936a>] common_interrupt+0x16a/0x16a
  
```

```

/usr/share/bcc/tools/trace 'r::sk_filter_trim_cap (retval != 0) "%d", retval'
281 281 ksoftirqd/54 sk_filter_trim_cap -12
281 281 ksoftirqd/54 sk_filter_trim_cap -12
20818 81613 bwa sk_filter_trim_cap -12
20818 81613 bwa sk_filter_trim_cap -12
314457 314457 kworker/54:1H sk_filter_trim_cap -12
314457 314457 kworker/54:1H sk_filter_trim_cap -12

/*
 * If the skb was allocated from pfmemalloc reserves, only
 * allow SOCK_MEMALLOC sockets to use it as this socket is
 * helping free memory
 */
if (skb_pfmemalloc(skb) && !sock_flag(sk, SOCK_MEMALLOC)) {
    NET_INC_STATS(sock_net(sk), LINUX_MIB_PFMEMALLOCDROP);
    return -ENOMEM;
}
  
```

# Lustre client 内存问题2

CentOS 7的vm.min\_free\_kbytes默认值太小

当达到内存最低水位线，则会导致响应时间过长，从而导致丢包。

```

Call Trace:
<IRQ>
[<fffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
[<fffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
<IRQ>
[<fffffffff88f7ff85>] dump_stack+0x19/0x1b
[<fffffffff88f7ff85>] dump_stack+0x19/0x1b
[<fffffffff88ec46ba>] ? tcp_v4_do_rcv+0x10a/0x350
[<fffffffff88ec46ba>] ? tcp_v4_do_rcv+0x10a/0x350
<IRQ> [<fffffffff88f7ff85>] dump_stack+0x19/0x1b
[<fffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
[<fffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
[<fffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffffff889c4ac0>] warn_alloc_failed+0x110/0x180
[<fffffffff88e81286>] ? sch_direct_xmit+0x86/0x250
[<fffffffff88e81286>] ? sch_direct_xmit+0x86/0x250
[<fffffffffc04d8036>] ? iptable_filter_hook+0x36/0x80 [iptable_filter]
[<fffffffffc04d8036>] ? iptable_filter_hook+0x36/0x80 [iptable_filter]
[<fffffffff88e6ed01>] ? sk_filter_trim_cap+0x31/0xe0
[<fffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450
[<fffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450
[<fffffffff88f7b4a0>] __alloc_pages_slowpath+0x6bb/0x729
[<fffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450
[<fffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450
[<fffffffff88a18e18>] __alloc_pages_current+0x98/0x110
[<fffffffff88a18e18>] __alloc_pages_current+0x98/0x110
[<fffffffffc04d8036>] ? iptable_filter_hook+0x36/0x80 [iptable_filter]
[<fffffffff88a18e18>] __alloc_pages_current+0x98/0x110
[<fffffffff88a18e18>] __alloc_pages_current+0x98/0x110
[<fffffffff889e5748>] kmalloc_order+0x18/0x40
[<fffffffff889e5748>] kmalloc_order+0x18/0x40
[<fffffffff889c9146>] __alloc_pages_nodemask+0x436/0x450

kswapd0: page allocation failure: order:2, mode:0x4020
kswapd0: page allocation failure: order:2, mode:0x4020
kswapd0: page allocation failure: order:2, mode:0x4020

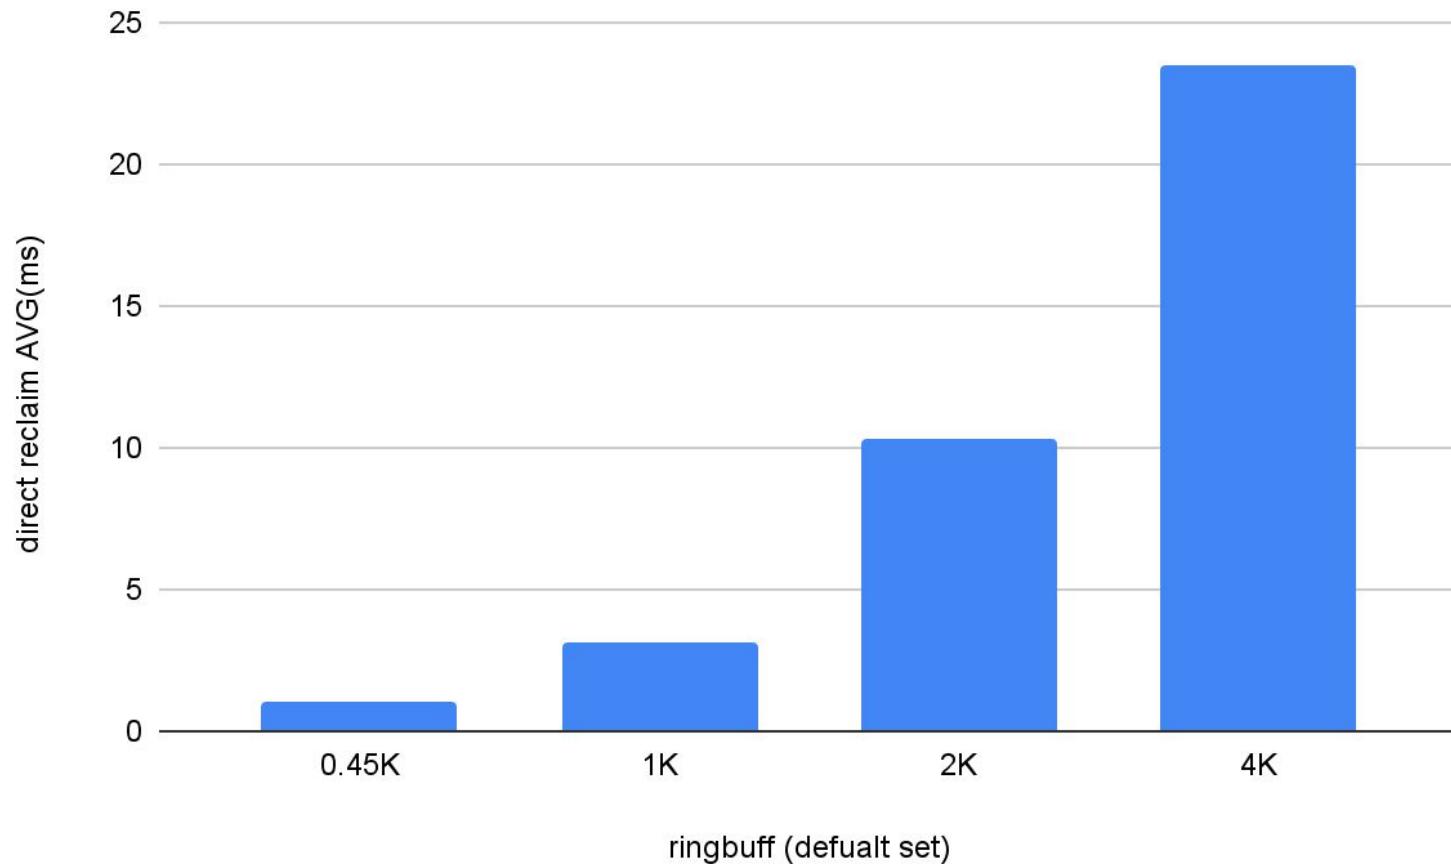
#define GFP_HIGH 0x20u
#define GFP_COMP 0x4000u

```

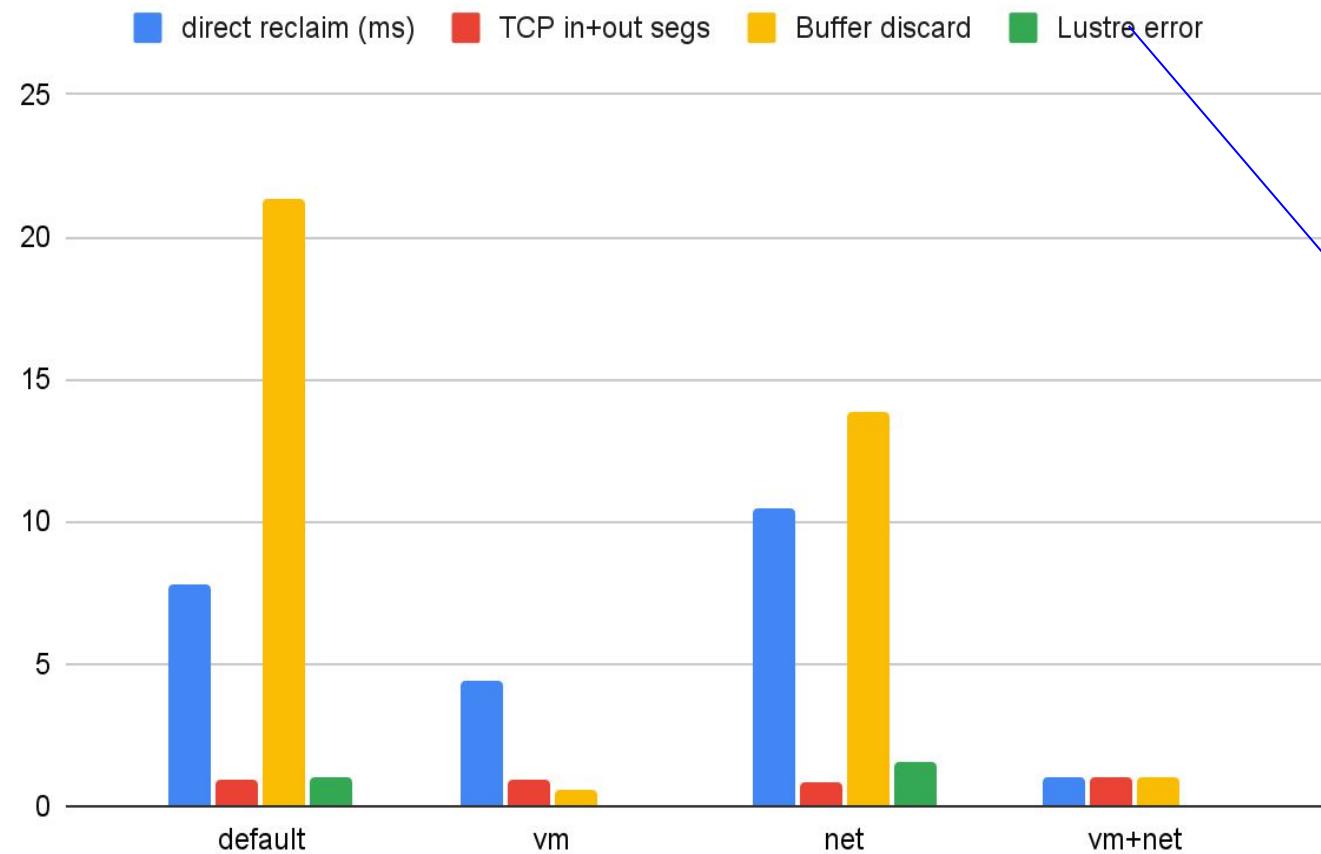
# 测试

- Loop test
  - The user application malloc/copy/release 90% memory
  - 48 x copy in lustre
  - 16 x network test application
  - No trigger swap
- Hardware
  - 2 x E5 2643v3
  - 8x 16G DDR 2133MHz
  - CentOS 7.8(3.10.0-1127.el7.x86\_64)
  - 9000 MTU
  - 2x 10GbE SR SFP+
- Test Parameters
  - Mem
    - Min\_free\_kbytes 700000
    - Zone\_reclaim\_mode 1
  - Net
    - Base\_reachable\_time\_ms 1200000
    - Netdev\_max\_backlog 4096
    - Somaxconn 4096
    - Dev\_weight 128
    - Netdev\_budget 600
  - NIC ringbuff
    - 453, 1000, 2000, 4078
  - Lustre
    - Server and client: 2.12.6, 2x MDS, 6x OSS
    - max\_dirty\_mb=2000
    - max\_read\_ahead\_mb=1024
    - grant\_shrink=0
    - max\_pages\_per\_rpc=256
    - osc max\_rpccs\_in\_flight=8

# Direct reclaim latency with ring buffer



# 比较

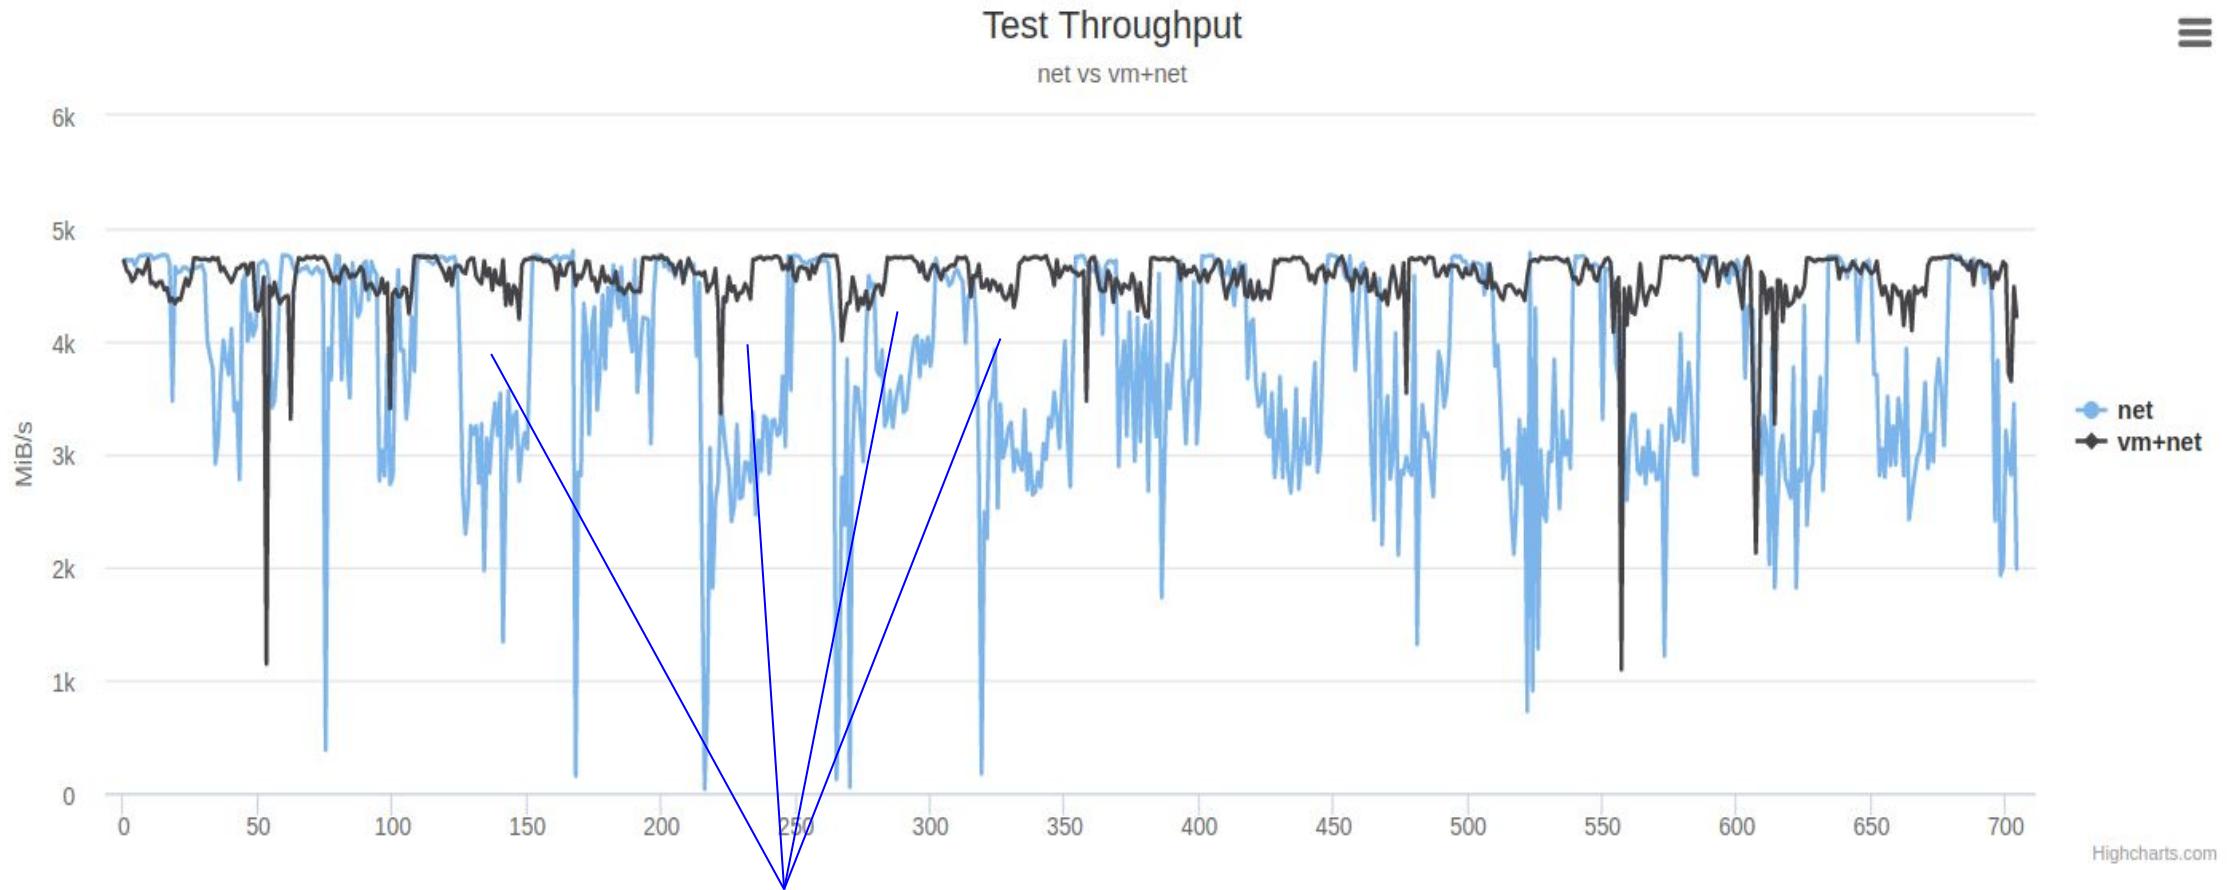


- Latency times
  - The lower the better
- TCP in+out segs times
  - The higher the better
- buff discard times
  - The lower the better
- Lustre Errors

Network err  
Connection lost  
Connection restore  
Page discard

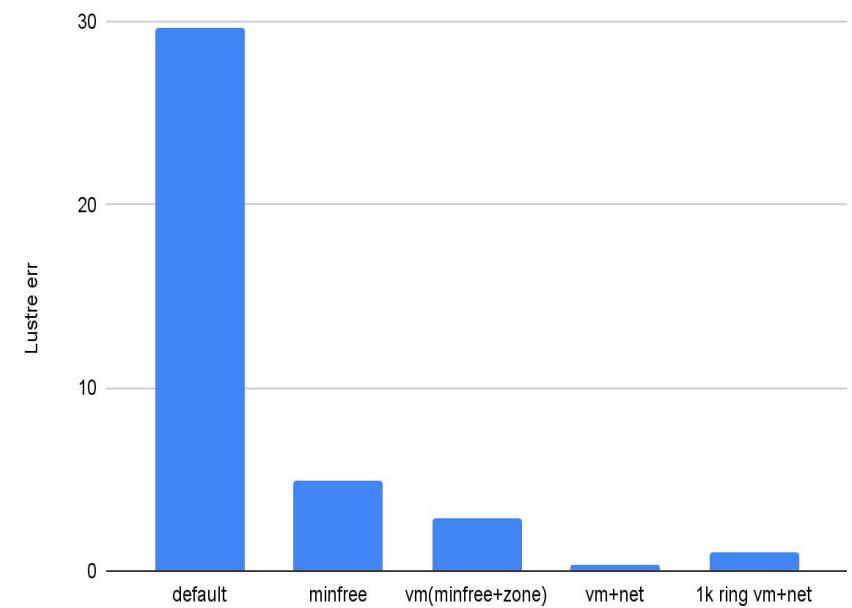
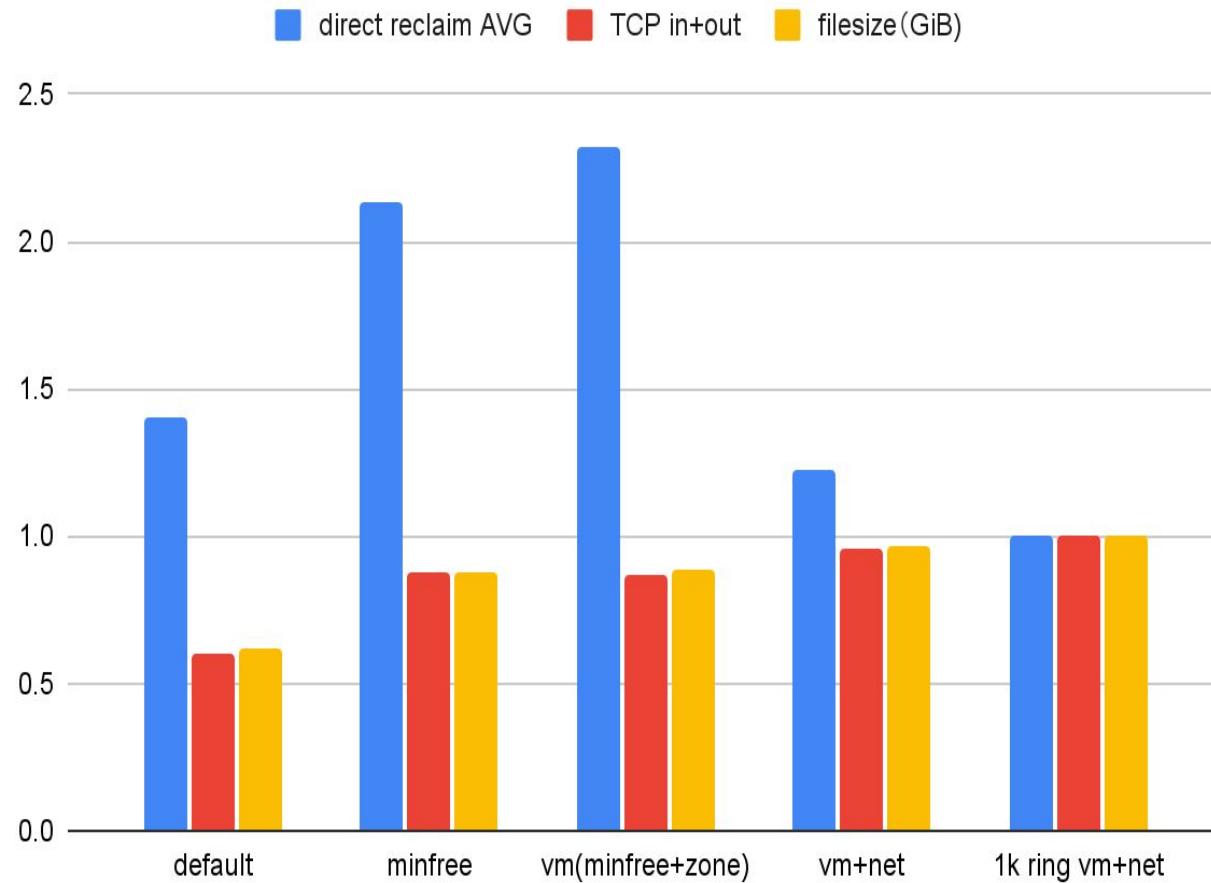
Cannot send after transport endpoint shutdown  
input/output error

# 重载下的调整效果



Total throughput, the higher the better, vm parameter is important

# 更激烈的内存竞争



# 多方位的追踪丢包

106x OSS

468x compute 相对较多的数据传输

发现了301 compute 出现丢包

107 x 25GbE:

TX: 102397230 RX: 102395445

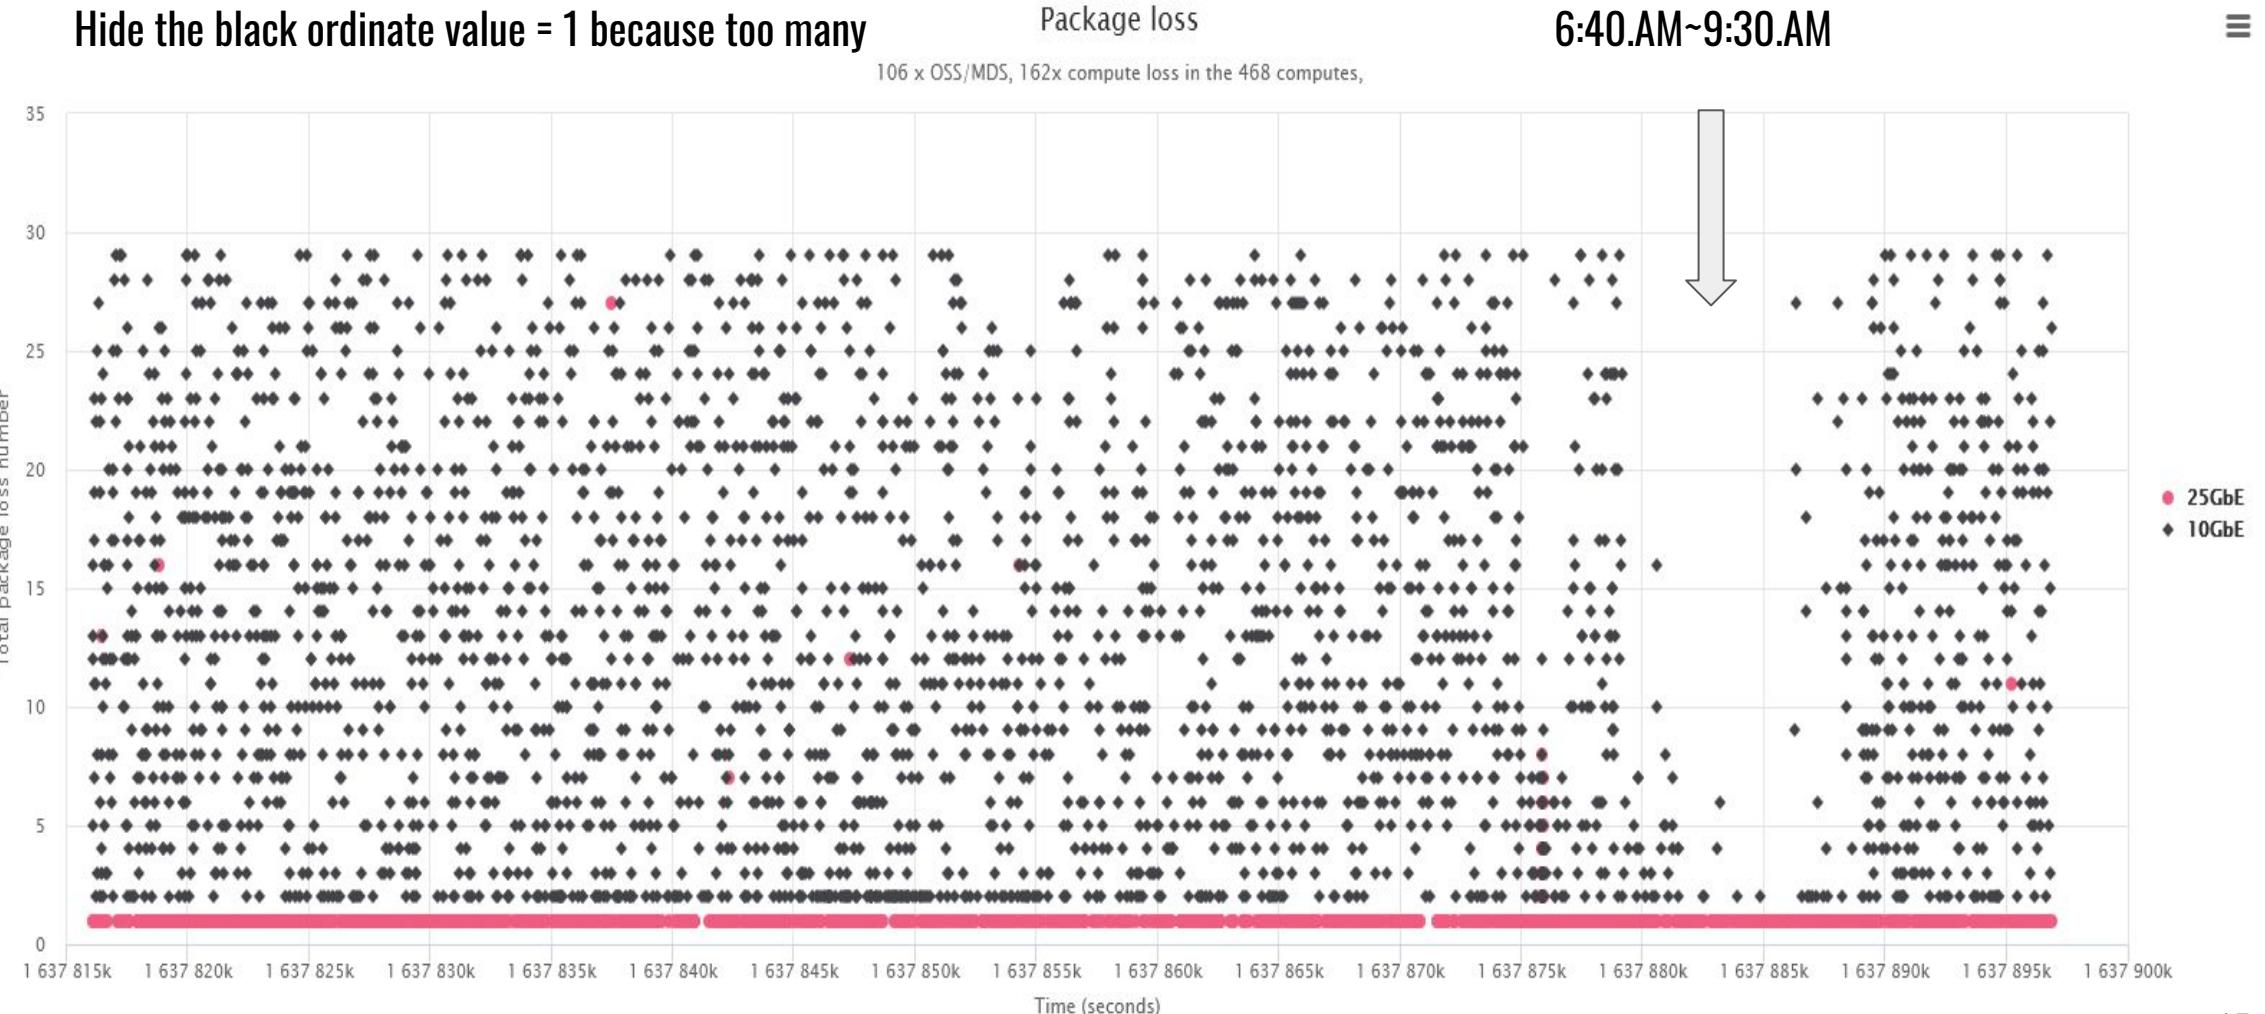
Ratio: 0.001743%

175 x 10GbE:

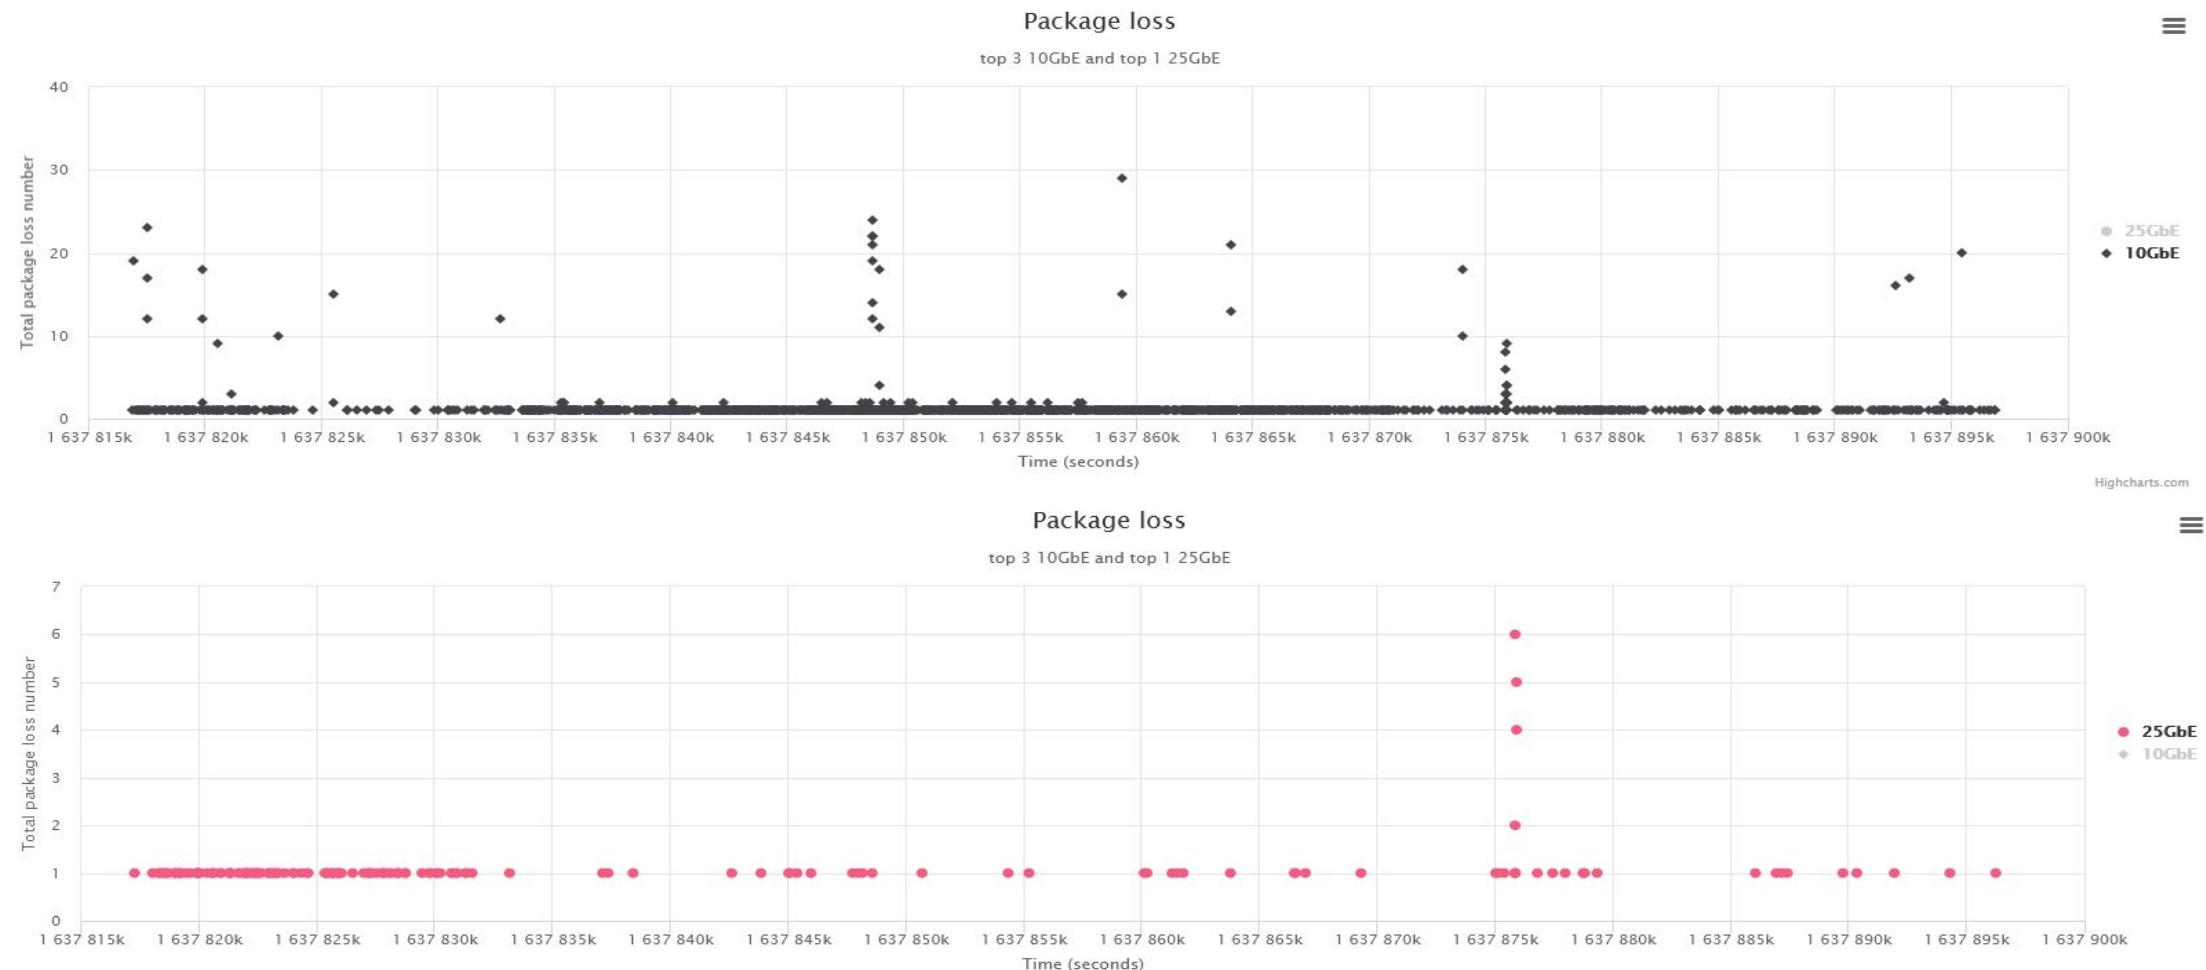
TX: 110814090 RX: 110760277

Ratio: 0.048562%

# 低负载下比较两种端口丢包总量



# 低负载节点单机比较(CPU/mem<40%, NIC BW<400MB/s)



# Tips

# BMC Watchdog

ipmi watchdog service does not support it

Countdown

```
-----  
|  
# ipmitool raw 0x06 0x24 0x44 0x20 0x1 0x10  
0x1 0x48  
# ipmitool mc watchdog reset  
|  
----if watchdog reset failed, trigger power  
diag
```

Use the high priority process

```
# ipmitool mc watchdog get  
Watchdog Timer Use: SMS/OS (0x44)  
Watchdog Timer Is: Started/Running  
Watchdog Timer Actions: No action (0x20)  
Pre-timeout interval: 1 seconds  
Timer Expiration Flags: 0x00  
Initial Countdown: 1843 sec  
Present Countdown: 1837 sec ←countdown
```

kernel.panic\_on\_unrecovered\_nmi=1  
Power diag trigger the kdump

```
[ 6.374483] 7010 total pagecache pages  
[ 6.375152] 0 pages in swap cache  
[ 6.375899] Swap cache stats: add 0, delete 0, find 0/0  
[ 6.376651] Free swap = 0kB  
[ 6.377418] Total swap = 0kB  
[ 6.378140] 43004 pages RAM  
[ 6.378857] 0 pages HighMem/MovableOnly  
[ 6.379570] 13888 pages reserved  
[ 6.380272] [ pid ] uid tgid total_vm rss nr_ptes swapents oom_score_adj name  
[ 6.380987] [ 78] 0 78 8747 681 21 0 0 systemd-journal  
[ 6.381785] [ 83] 0 83 8914 595 22 0 -1000 systemd-udevd  
[ 6.382428] [ 85] 0 85 11165 983 24 0 0 systemd-udevd  
[ 6.383132] [ 86] 0 86 11039 785 23 0 0 systemd-udevd  
[ 6.383842] [ 204] 0 204 3855 520 12 0 0 kdump.sh  
[ 6.384548] [ 236] 0 236 17473 4417 17 0 0 vmcore-dmesg  
[ 6.385252] Out of memory: Kill process 236 (vmcore-dmesg) score 147 or sacrifice child  
[ 6.385967] Killed process 236 (vmcore-dmesg), UID 0, total-vm:69892kB, anon-rss:17256kB, file-rss:4kB, shmem-rss:400kB  
/lib/kdump-lib/initramfs.sh: line 117: 236 Killed $_dmesg_collector /proc/vmcore > ${_path}/vmcore-dmesg-incomplete.txt  
kdump: saving vmcore-dmesg.txt failed  
kdump: saving vmcore  
Copying data : [ 27.0 %] \ eta: 1s[ 7.785835] random: crng init done  
Copying data : [100.0 %] \ eta: 0s  
kdump: saving vmcore complete  
[ 11.415722] systemd-shutdown[1]: Syncing filesystems and block devices.  
[ 11.416757] systemd-shutdown[1]: Sending SIGTERM to remaining processes...  
[ 11.420414] systemd-journal[78]: Received SIGTERM from PID 1 (systemd-shutdown).  
[ 11.421677] systemd-shutdown[1]: Sending SIGKILL to remaining processes...  
[ 11.424926] systemd-shutdown[1]: Unmounting file systems.  
[ 11.425798] systemd-shutdown[243]: Remounting '/sysroot' read-only in with options 'attr2,inode64,noquota'.  
[ 11.426986] systemd-shutdown[244]: Unmounting '/sysroot'.  
[ 11.434732] XFS (sda3): Unmounting Filesystem  
[ 11.435678] systemd-shutdown[245]: Remounting '/' read-only in with options 'size=51084k,nr_inodes=12826'.  
[ 11.436479] systemd-shutdown[246]: Remount  
[ 11.437213] systemd-shutdown[1]: All file  
[ 11.437867] systemd-shutdown[1]: Deactivat  
[ 11.438517] systemd-shutdown[1]: All swap  
[ 11.439128] systemd-shutdown[1]: Detachin  
[ 11.439784] systemd-shutdown[1]: All loop  
[ 11.440357] systemd-shutdown[1]: Detachin  
[ 11.440957] systemd-shutdown[1]: All DM d  
[ 11.441678] systemd-shutdown[1]: Syncing  
[ 11.442281] systemd-shutdown[1]: Rebootin  
Uhuh. NMI received for unknown reason 39 on CPU 0.  
Do you have a strange power saving mode enabled?  
Kernel panic - not syncing: NMI: Not continuing  
CPU: 0 PID: 0 Comm: swapper/0 Kdump: loaded Tainted: P 0E ----  
Hardware name:  
Call Trace:  
<NMI> [<fffffffffb317ff85>] dump_stack+0x19/0x1b  
[<fffffffffb3179521>] panic+0xe8/0x21f  
[<fffffffffb2a9ba6f>] nmi_panic+0x3f/0x40  
[<fffffffffb318adf6>] do_nmi+0x3f6/0x460  
[<fffffffffb3189d9c>] end_repeat_nmi+0x1e/0x81  
[<fffffffffb3188114>] ? intel_idle+0xd4/0x225  
[<fffffffffb3188114>] ? intel_idle+0xd4/0x225  
[<fffffffffb3188114>] ? intel_idle+0xd4/0x225  
<OE> [<fffffffffb2fc6015>] cpuidle_enter_state+0x45/0xd0  
[<fffffffffb2fc617e>] cpuidle_idle_call+0xde/0x230  
[<fffffffffb2a37c6e>] arch_cpu_idle+0xe/0xc0  
[<fffffffffb2b01c2a>] cpu_startup_entry+0x14a/0x1e0  
[<fffffffffb316e687>] rest_init+0x77/0x80  
[<fffffffffb378b1cf>] start_kernel+0x44b/0x46c  
[<fffffffffb378ab84>] ? repair_env_string+0x5c/0x5c  
[<fffffffffb378a120>] ? early_idt_handler_array+0x120/0x120  
[<fffffffffb378a738>] x86_64_start_reservations+0x24/0x26  
[<fffffffffb378a88e>] x86_64_start_kernel+0x154/0x177  
[<fffffffffb2a000d5>] start_cpu+0x5/0x14
```

# Uncorrected errors

Gigabytes processed [10^9 bytes]	Total uncorrected errors
518256.725	8
70731.604	0
100.017	0

厂商的日志系统并不告警  
脚本通过smartctl来扫描

# Udev帮助进一步规避brain split

```
ACTION=="add|change", KERNEL=="sd*", RUN+="/bin/sh -c 'echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6015 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6051 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6021 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f0f1 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f0e5 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f359 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f121 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6045 | cut -d/ -f7)/device/delete; \  
'"  
ACTION=="add|change", KERNEL=="sd*", RUN+="/bin/sh -c 'echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6009 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f0f5 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f0e9 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219b7d79 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee2199f0d5 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f5fdd | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f6005 | cut -d/ -f7)/device/delete; \  
echo 1 > /sys/block/$(ls -l /dev/disk/by-id/scsi-358ce38ee219f109 | cut -d/ -f7)/device/delete; \  
'"
```

和MMP一起构成双保险

# 谢谢

# THANKS

Owned by All, Completed by All and Shared by All

共有 · 共为 · 共享