Performance

Comparison of DragonFlyBSD and other OSes.

- DragonFlyBSD.
  - 719bf70a37139bc3bedc84ab0975df7107155714.
  - Label: dragonfly.
  - Label: dragonfly-invariant, kernel with “options INVARIANTS”.
- FreeBSD.
  - Label: freebsd, r314268, GENERIC-NODEBUG.
- Linux.
  - Label: linux3.10, Linux centos7 3.10.0-514.6.2.el7.x86_64.
  - Label: linux4.9, Linux deb 4.9.0-1-amd64, with 24 RX/TX rings.
  - Label: linux4.9-16rxring, Linux deb 4.9.0-1-amd64, with 16 RX/TX rings.

Performance HTTP/1.1 short lived connections

Configuration:

- Nginx ‘reuseport’ is enabled on DragonFlyBSD and Linux.
  - Not applicable to FreeBSD.
- DragonFlyBSD polling(4) @1000Hz.
- Linux tuning.
  - Firewall is disabled.
  - SELinux is disabled.
  - NOFILE is increased to 40960 both system-wise and for each nginx worker.
  - net.core.netdev_max_backlog=65535
  - net.core.somaxconn=256

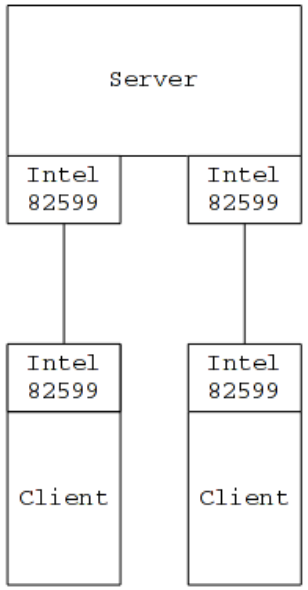

Server:
2x E5-2620v2, 32GB DDR3-1600, Hyperthread enabled.

nginx:
Installed from dports, version 1.11.8

nginx.conf:
Access log is disabled. 32 workers, 16K connections/worker.

15K concurrent connections on each client.

Client:
i7-3770, 16GB DDR3-1600, Hyperthread enabled.

MSL on clients and server are changed to 10ms by:
route change -net net -msl 10

/boot/loader.conf:
kern.ipc.nmbclusters=524288

/etc/sysctl.conf:
machdep.mwait.CX.idle=AUTODEEP
kern.ipc.somaxconn=256
net.inet.ip.portrange.last=40000

Figure 1.15

Performance
HTTP/1.1 short lived connections

HTTP/1.1 traffic generator:

- https://github.com/sepherosa/wrk
  - sephe/wrk branch.
- Each client ran:

```
wrk -c 15000 -t 8 -d 120s --delay --latency --connreqs 1 http://server/X_K.bin
```

- Each client can generate 160Krequests/s.
- 1K.bin, 8K.bin, and 16K.bin are static files.
  - DragonFlyBSD on hammerfs.
  - Linux on ext4.
  - FreeBSD on UFS2.

Performance
HTTP/1.1 short lived connections

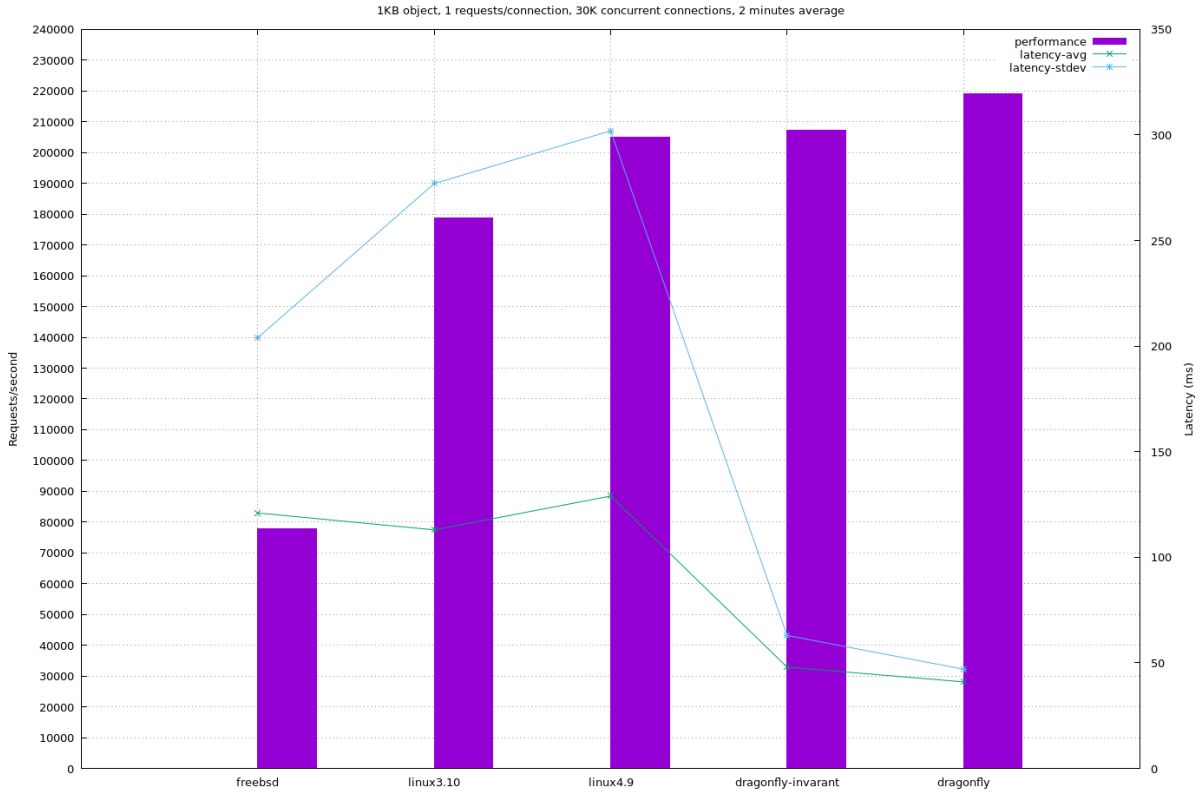

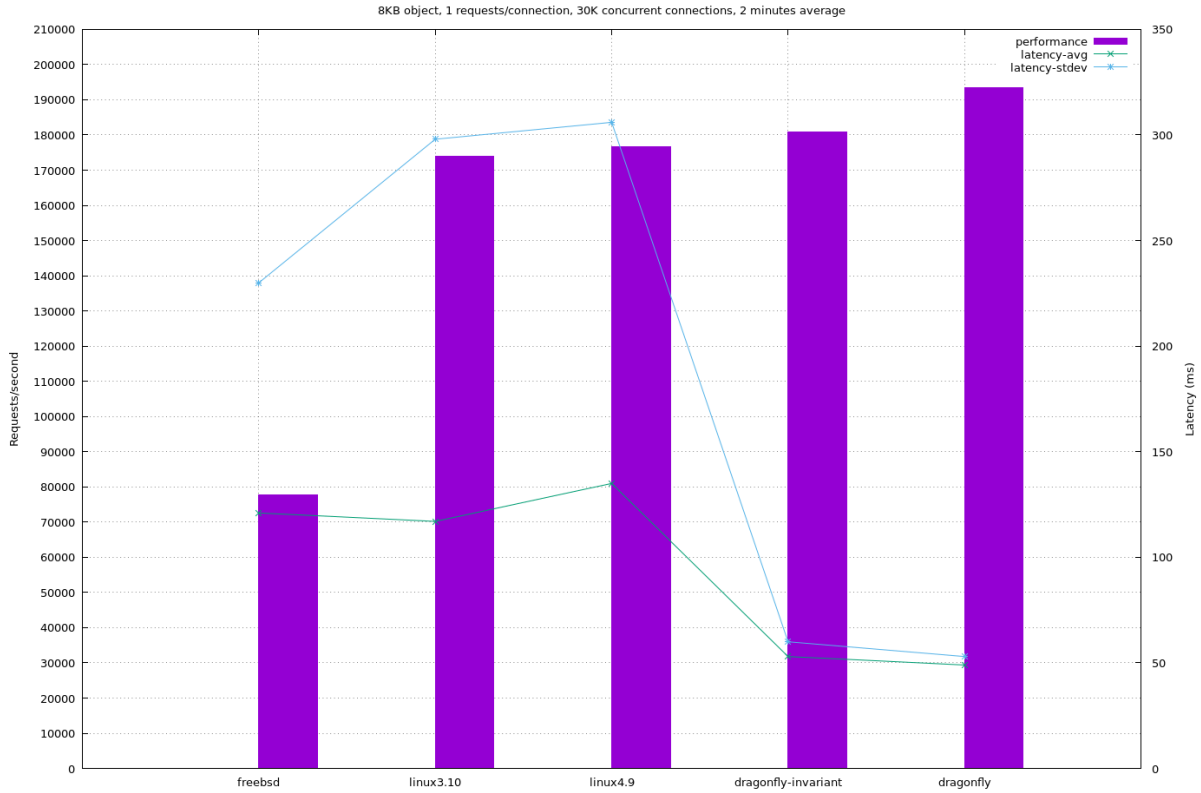

Performance
HTTP/1.1 short lived connections

DragonFlyBSD and Linux 4.9 reach 9.5Gbps on each NIC with 16KB web object.

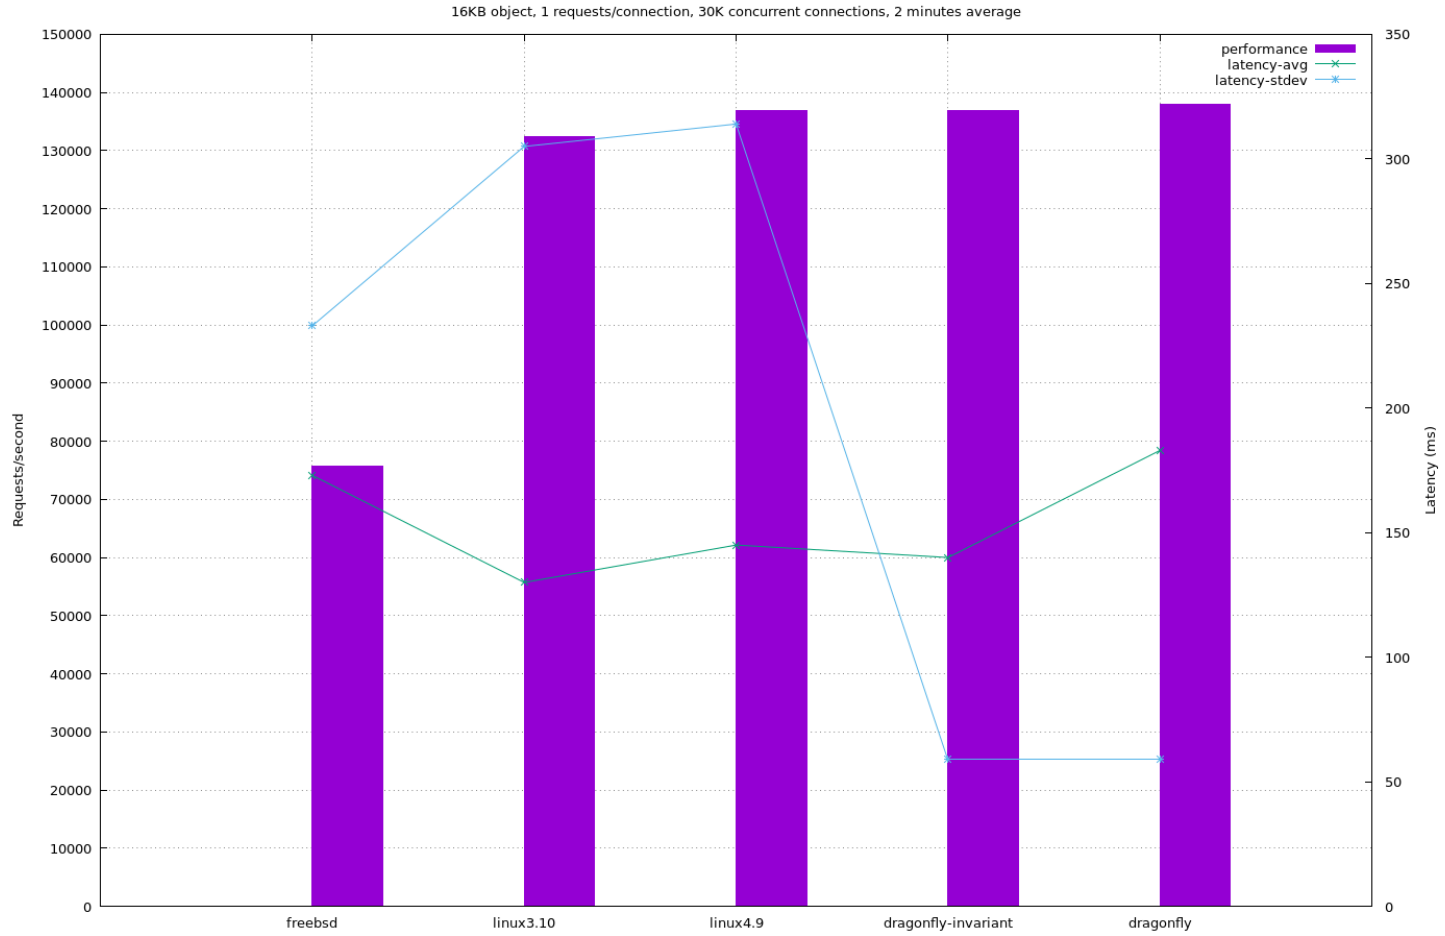

Performance
Bi-directional IPv4 forwarding

Configuration:

- DragonFlyBSD polling(4).
  - 6000Hz, default polling rate.
  - Reserved 5% cpu time for userspace.
- Linux tuning.
  - Firewall is disabled.
  - SELinux is disabled.
  - net.core.netdev_max_backlog=65535

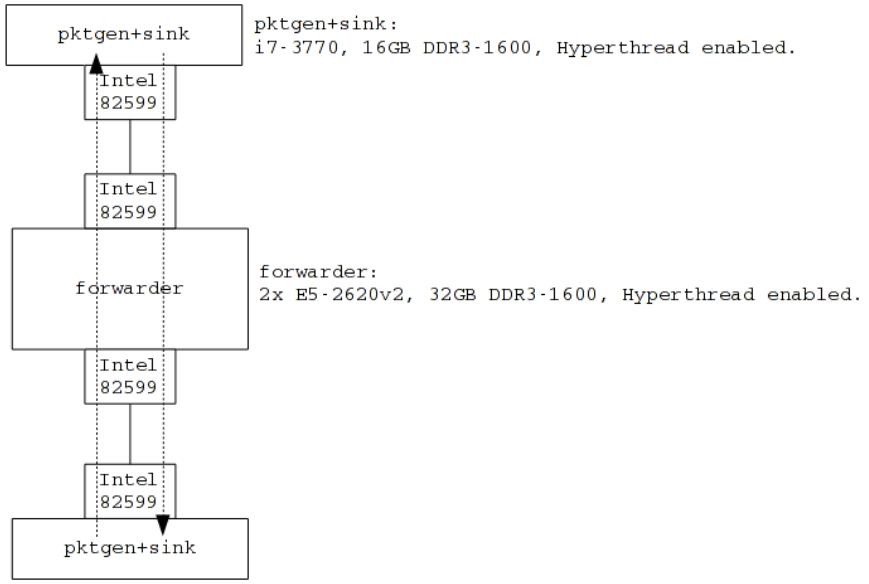


```
/boot/loader.conf:
kern.ipc.nmbclusters=5242880

/etc/sysctl.conf:
machdep.mwait.CX.idle=AUTODEEP

Traffic generator:
DragonFlyBSD's in kernel packet generator, which
has no issue to generate 14.8Mpps.

Traffic sink:
ifconfig ix0 monitor

Flow control is disabled on all testing machines:
ifconfig ix0 -mediaopt flowcontrol

Traffic:
Each pktgen targets 208 IPv4 addresses, which are
mapped to one link layer address on 'forwarder'.

Performance counting:
Only packets sent by 'forwarder' are counted.
```

Figure 1.16

Performance Bi-directional IPv4 forwarding

- DragonFlyBSD is impacted quite a lot by ‘options INVARIANTS’ for this test.
- Linux 4.9 Intel 82599 driver uses “flow director”.
  - Can use all 64 RX/TX rings.
  - 24 RX/TX rings are used here.
  - Not RSS, i.e. DragonFlyBSD cannot use it.
- Just for comparison purpose, Linux 4.9 using 16 RX/TX rings (linux4.9-16rxring) is benched.
  - Slower than DragonFlyBSD with the same number of RX/TX rings.