

# A study on ARM vs Intel networking performance

Focused on power consumption

Marcelo Ricardo Leitner  
SrPr Sw Eng  
mleitner@redhat.com

# Why?

- ▶ ARM is advertised as more power efficient. Is that really so?
- ▶ Customers are considering these new servers. But how does it scale?

# This study includes:

- ▶ TCP and UDP throughput test, IPv4
- ▶ CPU usage
- ▶ System power consumption
- ▶ Initial analysis

# Lab details

Three different systems:

- ▶ Ampere Altra ARMv8 (Q80-30) 3.00GHz Gigabyte R272-P31-00
- ▶ Intel(R) Xeon(R) CPU E5-2643 v3 @ 3.40GHz
- ▶ 2x Intel(R) Xeon(R) CPU E5-2690 v4 @ 2.60GHz
- ▶ All of them with Connex6-Dx @ 100Gbps
- ▶ All of them running 5.14.0-366.el9

## Config details

- ▶ NIC IRQs and iperf3 are pinned to dedicated CPUs
- ▶ Such cores are isolated (isolcpus)
- ▶ P-states adjusted to always have max freq, but left C-states on (tuned throughput-performance profile)

# Test procedure

- ▶ TCP throughput stream
- ▶ UDP throughput tests for small and MTU sized (1500) packets
- ▶ No odd iperf3 params
- ▶ With all combinations of the test systems

## Detailed results sample

```

test      tcp-tput
host      netdev19 netdev89
role      srv(rcv)  clt(snd)
tput Gbits/sec 31.8    31.8
power (W)  105     198
idle (W)   83      180
CPU2 %usr  6.28    0.80
CPU2 %sys  71.24   43.14
CPU2 %irq  0.32    0.21
CPU2 %soft 0.00    0.01
CPU2 %idle 22.16   55.85
CPU4 %usr  0.00    0.00
CPU4 %sys  0.15    0.11
CPU4 %irq  1.42    1.58
CPU4 %soft 70.20   27.38
CPU4 %idle 28.23   70.93
%CPU used  149.61  73.22
CPU/tput  4.704   2.302
W/tput    .691    .566

```

```

test      tcp-tput
host      netdev19 arm
role      srv(rcv)  clt(snd)
tput Gbits/sec 34.9    34.9
power (W)  105     149
idle (W)   83      128
CPU2 %usr  6.41    0.45
CPU2 %sys  73.51   42.39
CPU2 %irq  0.34    0.00
CPU2 %soft 0.00    0.00
CPU2 %idle 19.74   57.16
CPU4 %usr  0.00    0.00
CPU4 %sys  0.21    0.19
CPU4 %irq  1.57    0.00
CPU4 %soft 77.04   35.48
CPU4 %idle 21.18   64.33
%CPU used  159.08  78.51
CPU/tput  4.558   2.249
W/tput    .630    .601


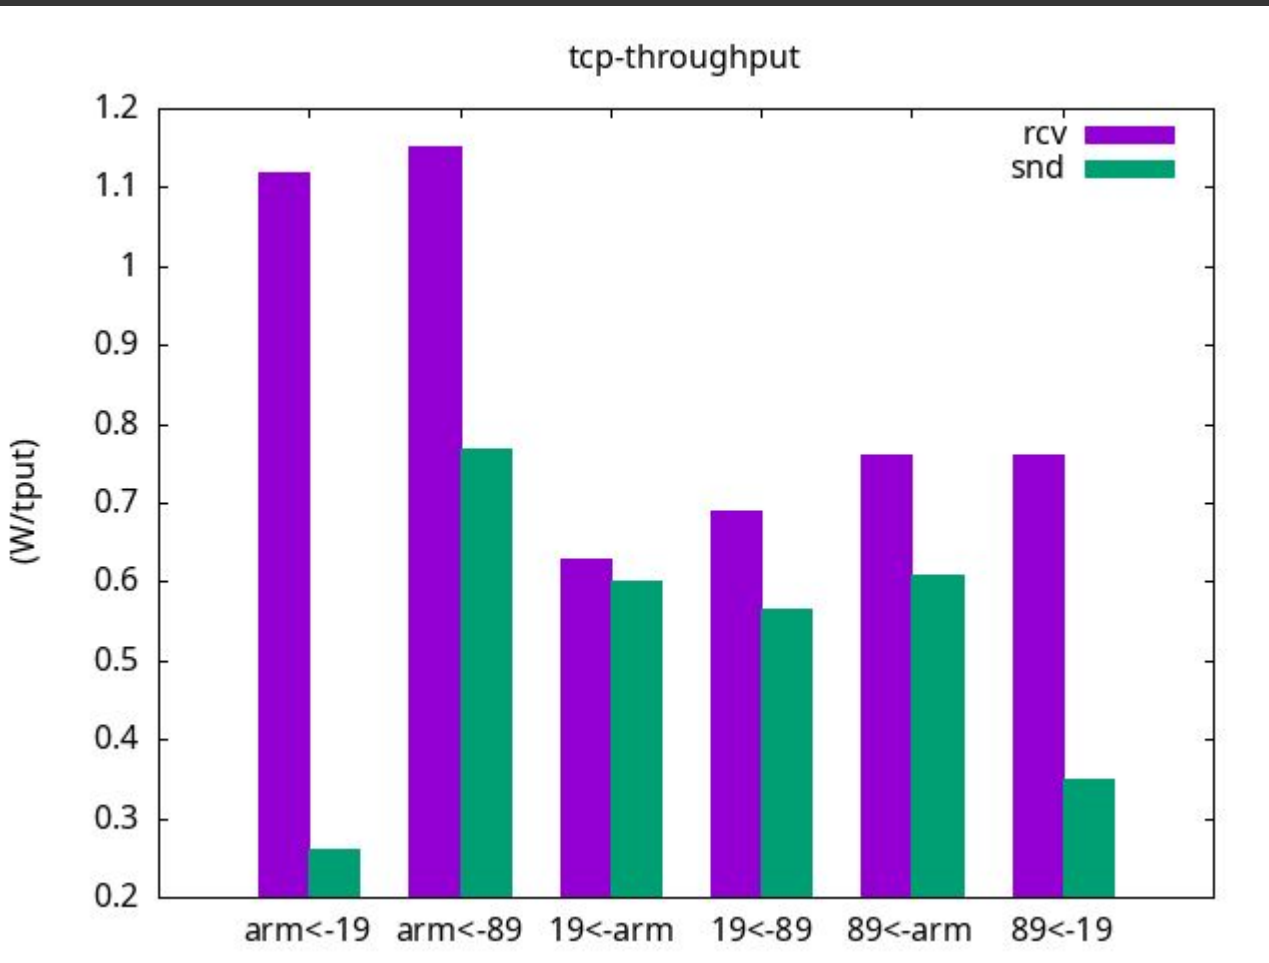
```

```

test      tcp-tput
host      arm      netdev19
role      srv(rcv)  clt(snd)
tput Gbits/sec 26.8    26.8
power (W)  158     90
idle (W)   128     83
CPU2 %usr  0.80    3.11
CPU2 %sys  98.02   74.57
CPU2 %irq  0.00    0.38
CPU2 %soft 0.00    0.00
CPU2 %idle 1.18    21.94
CPU4 %usr  0.00    0.00
CPU4 %sys  0.02    0.24
CPU4 %irq  0.00    3.72
CPU4 %soft 98.16   23.01
CPU4 %idle 1.81    73.03
%CPU used  197.01  105.03
CPU/tput  7.351   3.919
W/tput    1.119   .261

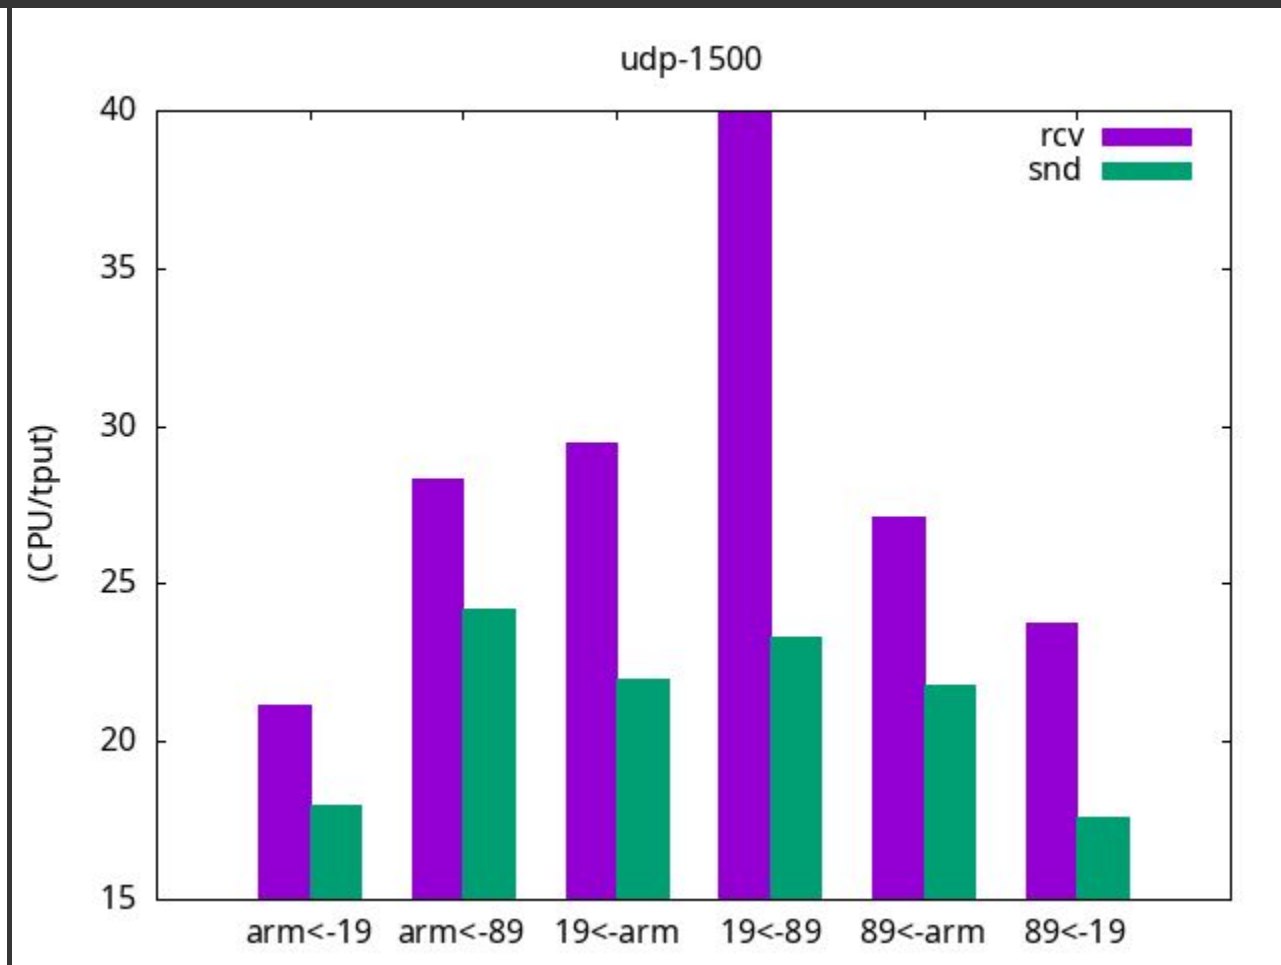

```

# TCP results summary (tput = Gbits/sec)

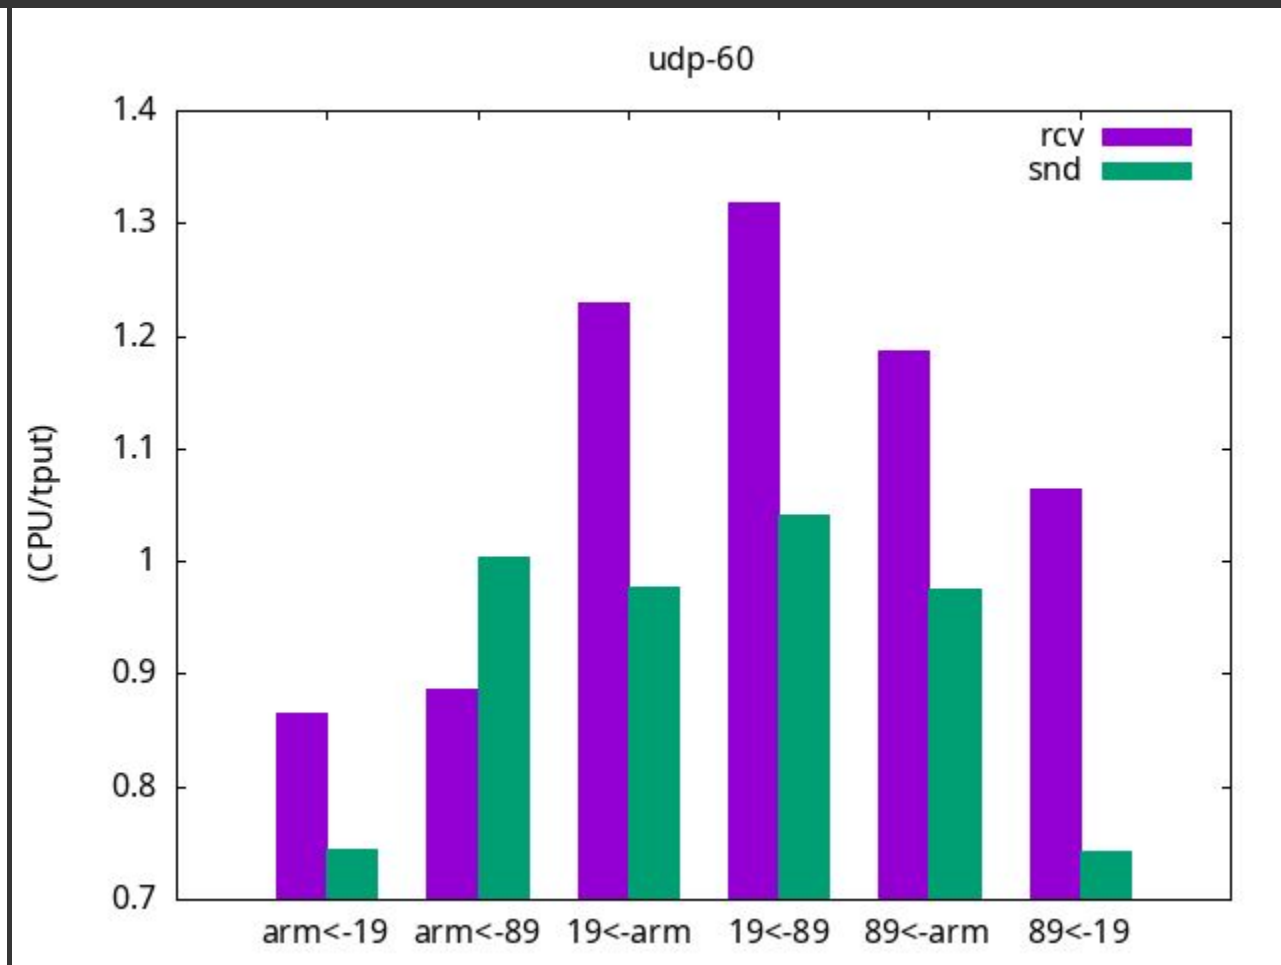
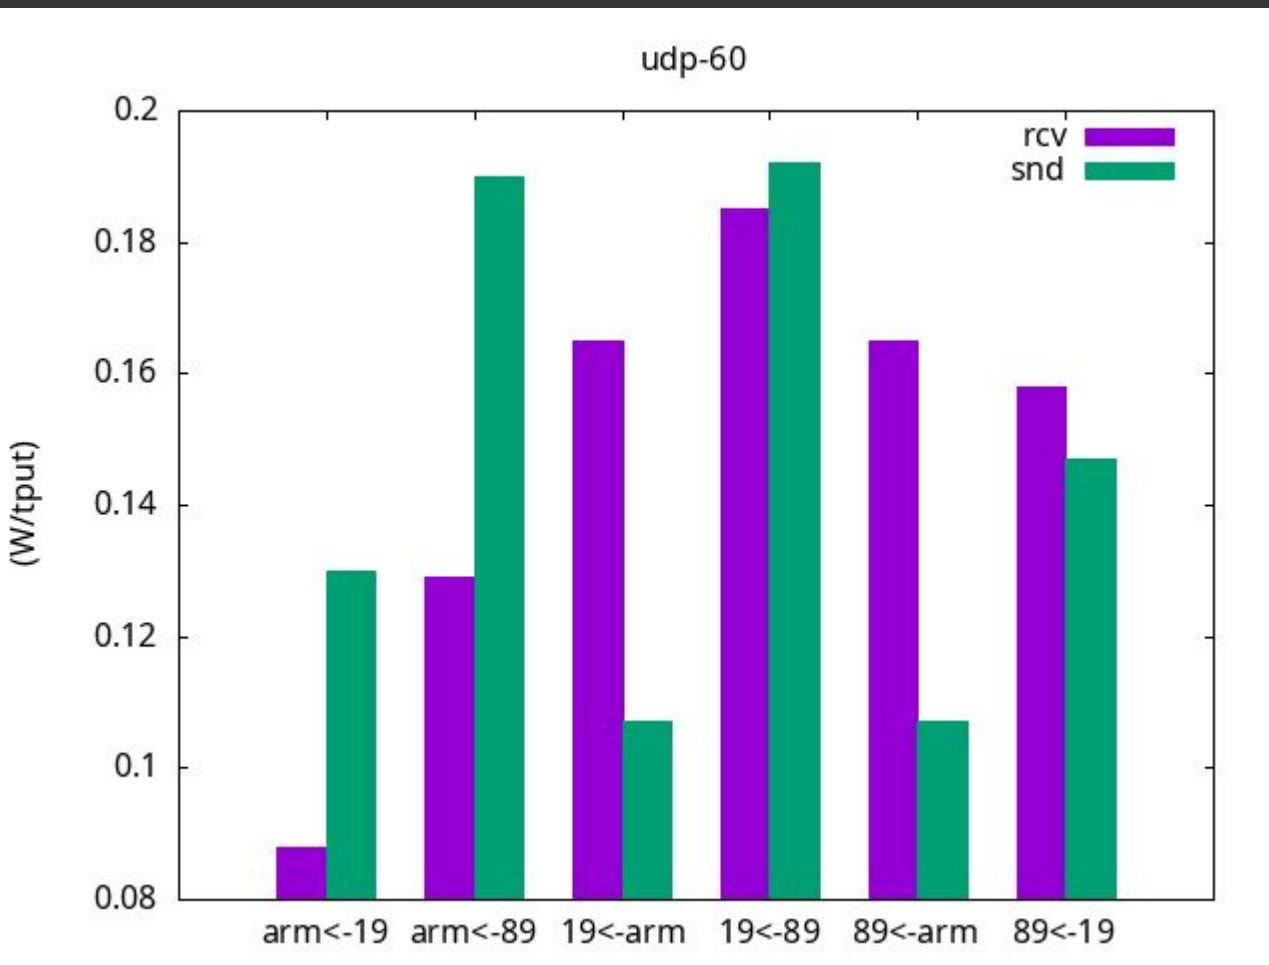




# UDP 1500 results summary (tput = Gbits/sec)



# UDP 60 summary (tput = Mbits/sec)



# ARM and E5-2643 v3, tcp rcv side

```
- 90.89% 0.72% iperf3 [kernel.kallsyms] [k] tcp_recvmsg_locked
- 90.17% tcp_recvmsg_locked
- 65.36% skb_copy_datagram_iter
- 65.15% __skb_datagram_iter
- 34.84% simple_copy_to_iter
  29.22% __arch_copy_to_user
- 4.45% __check_object_size
  - 4.33% __check_object_size.part.0
    - 3.80% check_heap_object
      1.62% memblock_is_map_memory
0.89% _copy_to_iter
- 28.42% __skb_datagram_iter
- 27.35% simple_copy_to_iter
  23.05% __arch_copy_to_user
- 3.45% __check_object_size
  - 3.36% __check_object_size.part.0
    - 2.96% check_heap_object
      1.15% memblock_is_map_memory
0.72% _copy_to_iter
- 18.71% __kfree_skb
- 18.28% skb_release_data
- 8.32% kfree_skb_list_reason
- 7.36% skb_release_data
- 2.87% __folio_put
  2.17% _raw_spin_unlock_irqrestore
  0.55% free_unref_page
- 0.56% skb_free_head
  page_frag_free
- 0.75% kfree_skbmem
  0.68% kmem_cache_free
- 4.24% __folio_put
  3.24% _raw_spin_unlock_irqrestore
  0.78% free_unref_page
```

```
- 69.57% 0.83% iperf3 [kernel.kallsyms] [k] tcp_recvmsg_locked
- 68.74% tcp_recvmsg_locked
- 52.89% skb_copy_datagram_iter
- __skb_datagram_iter
- 29.87% __skb_datagram_iter
- 18.69% _copy_to_iter
  15.96% copy_user_enhanced_fast_string
- 7.21% simple_copy_to_iter
  - 6.11% __check_object_size.part.0
    + 5.01% check_heap_object
- 13.54% _copy_to_iter
  11.58% copy_user_enhanced_fast_string
- 5.49% simple_copy_to_iter
  + 4.82% __check_object_size.part.0
- 7.17% __kfree_skb
- 7.12% skb_release_data
- 4.38% kfree_skb_list_reason
  3.37% skb_release_data
  0.54% __slab_free
+ 5.55% __tcp_transmit_skb
  1.10% __sk_mem_reduce_allocated
+ 0.84% __tcp_send_ack.part.0
0.29% 0.27% iperf3 libiperf.so.0.0.0 [.] iperf_tcp_rcv
```

## ARM tcp throughput rcv softirq side

```
16.70%  16.70%  ksoftirqd/4  [kernel.kallsyms]  [k] is_last_ethertype_ip.constprop.0
ret_from_fork
kthread
- smpboot_thread_fn
  - 16.70% run_ksoftirqd
    - 16.68% __do_softirq
      net_rx_action
      __napi_poll
      mlx5e_napi_poll
      mlx5e_poll_rx_cq
      mlx5e_rx_cq_process_basic_cqe_comp
    - mlx5e_handle_rx_cqe_mpwrq
      - 16.66% mlx5e_build_rx_skb
        is_last_ethertype_ip.constprop.0
```

# Conclusion

- ▶ Measuring power consumption for a workload is not easy
  - I.e., how to deal with power supply efficiency change?
- ▶ TCP receive seems weird and needs more digging
- ▶ TCP send and UDP send and receive seems to be okay

# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



[twitter.com/RedHat](https://twitter.com/RedHat)

# Detailed results

```

test          tcp-tput
host          arm          netdev19
role          srv(rcv)    clt(snd)
tput Gbits/sec 26.8        26.8
power (W)     158         90
idle (W)      128         83
CPU2 %usr     0.80        3.11
CPU2 %sys     98.02       74.57
CPU2 %irq     0.00        0.38
CPU2 %soft    0.00        0.00
CPU2 %idle    1.18        21.94
CPU4 %usr     0.00        0.00
CPU4 %sys     0.02        0.24
CPU4 %irq     0.00        3.72
CPU4 %soft    98.16       23.01
CPU4 %idle    1.81        73.03
%CPU used     197.01      105.03
CPU/tput      7.351       3.919
W/tput        1.119       .261

```

```

test          tcp-tput
host          arm          netdev89
role          srv(rcv)    clt(snd)
tput Gbits/sec 26.0        26.0
power (W)     158         200
idle (W)      128         180
CPU2 %usr     0.61        0.67
CPU2 %sys     94.72       31.73
CPU2 %irq     0.00        0.17
CPU2 %soft    0.00        0.00
CPU2 %idle    4.66        67.43
CPU4 %usr     0.00        0.00
CPU4 %sys     0.05        0.09
CPU4 %irq     0.00        1.68
CPU4 %soft    87.97       20.12
CPU4 %idle    11.98       78.11
%CPU used     183.36      54.46
CPU/tput      7.052       2.094
W/tput        1.153       .769

```

# Detailed results

```
test          udp-1500
host          arm          netdev19
role          srv(rcv)    clt(snd)
tput Gbits/sec 5.97        7.11
power (W)     145         105
idle (W)      128         83
CPU2 %usr     10.53       24.31
CPU2 %sys     70.39       75.44
CPU2 %irq     0.00        0.13
CPU2 %soft    0.00        0.00
CPU2 %idle    19.08       0.12
CPU4 %usr     0.00        0.00
CPU4 %sys     0.01        0.00
CPU4 %irq     0.00        1.45
CPU4 %soft    45.28       26.54
CPU4 %idle    54.71       72.00
%CPU used     126.21      127.88
CPU/tput      21.140      17.985
W/tput        2.847       3.094
```

```
test          udp-1500
host          arm          netdev89
role          srv(rcv)    clt(snd)
tput Gbits/sec 4.39        4.76
power (W)     144         202
idle (W)      128         180
CPU2 %usr     8.73        7.20
CPU2 %sys     63.70       92.65
CPU2 %irq     0.00        0.14
CPU2 %soft    0.00        0.00
CPU2 %idle    27.57       0.01
CPU4 %usr     0.00        0.00
CPU4 %sys     0.01        0.00
CPU4 %irq     0.00        1.98
CPU4 %soft    51.94       13.26
CPU4 %idle    48.05       84.76
%CPU used     124.38      115.23
CPU/tput      28.332      24.207
W/tput        3.644       4.621
```



# Detailed results

```

test          udp-60
host          arm          netdev19
role          srv(rcv)  clt(snd)
tput Mbits/sec 169          169
power (W)     143          105
idle (W)      128          83
CPU2 %usr     14.87        26.67
CPU2 %sys     82.46        73.19
CPU2 %irq     0.00         0.14
CPU2 %soft    0.00         0.00
CPU2 %idle    2.66         0.01
CPU4 %usr     0.00         0.00
CPU4 %sys     0.00         0.00
CPU4 %irq     0.00         0.65
CPU4 %soft    49.09        25.22
CPU4 %idle    50.91        74.13
%CPU used     146.43       125.86
CPU/tput      .866         .744
W/tput        .088         .130

```

```

test          udp-60
host          arm          netdev89
role          srv(rcv)  clt(snd)
tput Mbits/sec 108          110
power (W)     142          201
idle (W)      128          180
CPU2 %usr     8.66         7.55
CPU2 %sys     53.29        92.30
CPU2 %irq     0.00         0.15
CPU2 %soft    0.00         0.00
CPU2 %idle    38.05        0.01
CPU4 %usr     0.00         0.00
CPU4 %sys     0.00         0.00
CPU4 %irq     0.00         0.34
CPU4 %soft    33.87        10.10
CPU4 %idle    66.13        89.55
%CPU used     95.82        110.44
CPU/tput      .887         1.004
W/tput        .129         .190

```

# Detailed results

```

test          tcp-tput
host          netdev19  arm
role          srv(rcv)  clt(snd)
tput Gbits/sec 34.9      34.9
power (W)     105       149
idle (W)      83        128
CPU2 %usr     6.41      0.45
CPU2 %sys     73.51     42.39
CPU2 %irq     0.34      0.00
CPU2 %soft    0.00      0.00
CPU2 %idle    19.74     57.16
CPU4 %usr     0.00      0.00
CPU4 %sys     0.21      0.19
CPU4 %irq     1.57      0.00
CPU4 %soft    77.04     35.48
CPU4 %idle    21.18     64.33
%CPU used     159.08    78.51
CPU/tput      4.558     2.249
W/tput        .630      .601

```

```

test          tcp-tput
host          netdev19  netdev89
role          srv(rcv)  clt(snd)
tput Gbits/sec 31.8      31.8
power (W)     105       198
idle (W)      83        180
CPU2 %usr     6.28      0.80
CPU2 %sys     71.24     43.14
CPU2 %irq     0.32      0.21
CPU2 %soft    0.00      0.01
CPU2 %idle    22.16     55.85
CPU4 %usr     0.00      0.00
CPU4 %sys     0.15      0.11
CPU4 %irq     1.42      1.58
CPU4 %soft    70.20     27.38
CPU4 %idle    28.23     70.93
%CPU used     149.61    73.22
CPU/tput      4.704     2.302
W/tput        .691      .566

```

# Detailed results

```

test          udp-1500
host          netdev19  arm
role         srv(rcv)  clt(snd)
tput Gbits/sec 5.52      6.12
power (W)    106       143
idle (W)     83        128
CPU2 %usr    31.71     6.07
CPU2 %sys    56.70     92.68
CPU2 %irq    0.15      0.00
CPU2 %soft   0.00      0.00
CPU2 %idle   11.44     1.24
CPU4 %usr    0.00      0.00
CPU4 %sys    0.00      0.00
CPU4 %irq    0.25      0.00
CPU4 %soft   73.93     35.75
CPU4 %idle   25.82     64.25
%CPU used    162.74    134.51
CPU/tput     29.481    21.978
W/tput       4.166     2.450

```

```

test          udp-1500
host          netdev19  netdev89
role         srv(rcv)  clt(snd)
tput Gbits/sec 3.64      4.77
power (W)    93        201
idle (W)     83        180
CPU2 %usr    29.64     7.20
CPU2 %sys    52.10     92.64
CPU2 %irq    0.23      0.14
CPU2 %soft   0.00      0.00
CPU2 %idle   18.03     0.01
CPU4 %usr    0.00      0.00
CPU4 %sys    0.00      0.00
CPU4 %irq    0.51      0.68
CPU4 %soft   63.11     10.57
CPU4 %idle   36.38     88.75
%CPU used    145.59    111.24
CPU/tput     39.997    23.320
W/tput       2.747     4.402

```

## Detailed results

```

test          udp-60
host          netdev19  arm
role         srv(rcv)  clt(snd)
tput Mbits/sec 139      139
power (W)     106      143
idle (W)      83       128
CPU2 %usr     35.83    6.36
CPU2 %sys     60.91    93.63
CPU2 %irq     0.13     0.00
CPU2 %soft    0.00     0.00
CPU2 %idle    3.14     0.01
CPU4 %usr     0.00     0.00
CPU4 %sys     0.00     0.00
CPU4 %irq     0.23     0.00
CPU4 %soft    73.96    35.94
CPU4 %idle    25.81    64.06
%CPU used     171.05   135.93
CPU/tput      1.230    .977
W/tput        .165     .107

```

```

test          udp-60
host          netdev19  netdev89
role         srv(rcv)  clt(snd)
tput Mbits/sec 108      109
power (W)     103      201
idle (W)      83       180
CPU2 %usr     29.22    7.45
CPU2 %sys     50.29    92.41
CPU2 %irq     0.12     0.13
CPU2 %soft    0.00     0.00
CPU2 %idle    20.36    0.01
CPU4 %usr     0.00     0.00
CPU4 %sys     0.00     0.00
CPU4 %irq     0.48     1.36
CPU4 %soft    62.25    12.15
CPU4 %idle    37.27    86.49
%CPU used     142.37   113.50
CPU/tput      1.318    1.041
W/tput        .185     .192

```

## Detailed results

```

test      tcp-tput
host      netdev89  arm
role      srv(rcv)  clt(snd)
tput Gbits/sec  32.8      32.8
power (W) 205       148
idle (W)  180       128
CPU2 %usr  2.28      0.42
CPU2 %sys  75.11     38.30
CPU2 %irq  0.20      0.00
CPU2 %soft 0.01      0.00
CPU2 %idle 22.41     61.28
CPU4 %usr  0.00      0.00
CPU4 %sys  0.13      0.17
CPU4 %irq  1.30      0.00
CPU4 %soft 75.45     32.48
CPU4 %idle 23.12     67.35
%CPU used 154.47    71.37
CPU/tput  4.709     2.175
W/tput   .762      .609

```

```

test      tcp-tput
host      netdev89  netdev19
role      srv(rcv)  clt(snd)
tput Gbits/sec  34.2      34.2
power (W) 206       95
idle (W)  180       83
CPU2 %usr  2.48      2.78
CPU2 %sys  78.94     76.05
CPU2 %irq  0.22      0.37
CPU2 %soft 0.01      0.01
CPU2 %idle 18.35     20.79
CPU4 %usr  0.00      0.00
CPU4 %sys  0.22      0.47
CPU4 %irq  1.48      4.97
CPU4 %soft 77.32     33.29
CPU4 %idle 20.98     61.27
%CPU used 160.67    117.94
CPU/tput  4.697     3.448
W/tput   .760      .350

```

# Detailed results

```

test          udp-1500
host          netdev89  arm
role         srv(rcv)  clt(snd)
tput Gbits/sec 5.69      6.16
power (W)    202       143
idle (W)     180       128
CPU2 %usr    16.46     6.08
CPU2 %sys    69.53     92.77
CPU2 %irq    0.16      0.00
CPU2 %soft   0.00      0.00
CPU2 %idle   13.84     1.15
CPU4 %usr    0.00      0.00
CPU4 %sys    0.00      0.00
CPU4 %irq    0.15      0.00
CPU4 %soft   67.95     35.40
CPU4 %idle   31.90     64.60
%CPU used    154.26    134.25
CPU/tput     27.110    21.793
W/tput       3.866     2.435

```

```

test          udp-1500
host          netdev89  netdev19
role         srv(rcv)  clt(snd)
tput Gbits/sec 6.31      7.17
power (W)    203       108
idle (W)     180       83
CPU2 %usr    17.48     24.65
CPU2 %sys    75.24     74.80
CPU2 %irq    0.17      0.13
CPU2 %soft   0.00      0.00
CPU2 %idle   7.11      0.42
CPU4 %usr    0.00      0.00
CPU4 %sys    0.00      0.00
CPU4 %irq    0.12      1.44
CPU4 %soft   56.76     25.35
CPU4 %idle   43.11     73.21
%CPU used    149.78    126.37
CPU/tput     23.736    17.624
W/tput       3.645     3.486

```

## Detailed results

```

test          udp-60
host          netdev89  arm
role          srv(rcv)  clt(snd)
tput Mbits/sec 139      139
power (W)     203      143
idle (W)      180      128
CPU2 %usr     17.75    6.37
CPU2 %sys     75.11    93.62
CPU2 %irq     0.14     0.00
CPU2 %soft    0.00     0.00
CPU2 %idle    7.00     0.01
CPU4 %usr     0.00     0.00
CPU4 %sys     0.00     0.00
CPU4 %irq     0.14     0.00
CPU4 %soft    71.81    35.70
CPU4 %idle    28.04    64.30
%CPU used     164.96    135.69
CPU/tput      1.186    .976
W/tput        .165     .107

```

```

test          udp-60
host          netdev89  netdev19
role          srv(rcv)  clt(snd)
tput Mbits/sec 151      169
power (W)     204      108
idle (W)      180      83
CPU2 %usr     19.19    26.50
CPU2 %sys     80.51    73.36
CPU2 %irq     0.14     0.14
CPU2 %soft    0.00     0.00
CPU2 %idle    0.16     0.01
CPU4 %usr     0.00     0.00
CPU4 %sys     0.00     0.00
CPU4 %irq     0.13     0.65
CPU4 %soft    61.00    24.97
CPU4 %idle    38.88    74.38
%CPU used     160.96    125.61
CPU/tput      1.065    .743
W/tput        .158     .147

```