



Benchmark Report

Benchmark: eflatency, sfnt-pingpong

System: **ACE 3120-RZ**

Family: **ACE 3100-RZ**

Additional Equipment: AMD Alveo X3522 NIC

| Release Date | Revision | Notes |
|--------------|----------|-----------------------------|
| 01-Sept-25 | 1.0 | First release |
| 13-Oct-25 | 1.1 | Corrected SUT model details |



Contents

| | |
|---|----|
| Introduction..... | 3 |
| System Under Test (SUT)..... | 3 |
| Tuning | 4 |
| Grub Parameters..... | 4 |
| Tuned..... | 4 |
| Test Parameters..... | 5 |
| Results..... | 7 |
| Summary of overclocked results:..... | 7 |
| Summary of stock results:..... | 8 |
| Summary comparison at typical byte sizes..... | 9 |
| 4 Bytes..... | 9 |
| 64 Bytes..... | 9 |
| 128 Bytes..... | 10 |
| 256 Bytes | 10 |
| Conclusion..... | 10 |



Introduction

Blackcore provide fast, reliable, overclocked, liquid cooled servers that deliver the ultimate competitive edge to our clients within the electronic trading industry.

Testing

We will run a set of simple tests using eflatency and sfnt-pingpong^[1] to measure typical network latency in a Blackcore overclocked server system when coupled with an AMD low latency network interface card (NIC).

We will also compare stock^[2] performance to Blackcore's overclocked performance.

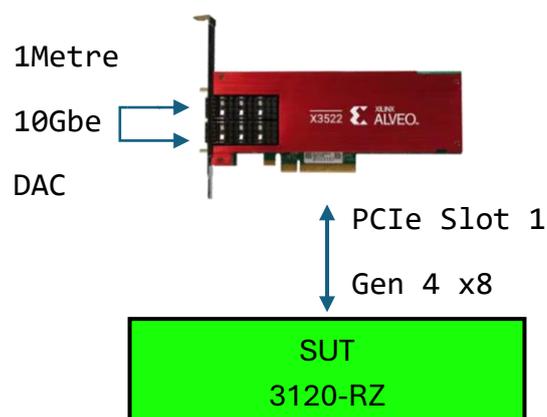
^[1] eflatency and sfnt-pinpong are provided by AMD as part of the Xilinx onload toolset.

See: [Xilinx-CNS/onload · GitHub](#)

^[2] Stock is defined as running the system without any BIOS tweaks, as the system would run "out of the box" and adhering to the CPU manufacturers default specifications. This includes Simultaneous Multi-Threading (SMT) / Hyperthreading (HT) and E-Cores enabled (where applicable). Stock testing is still performed in a Blackcore system and so may include cooling enhancements not found in other typical systems. Linux tuning will be identical to overclocked testing except for any adjustments needed to enable stock features (i.e.: enablement of SMT / HT).

System Under Test (SUT)

The SUT contains an AMD Alveo™ X3522 card in PCIe Slot 1 with a 1 Metre 10Gbe Direct Attach Cable (DAC) connected in loopback.





| Hardware | | Firmware |
|-------------------|--|--------------------------------------|
| Blackcore 3120-RZ | CPU: AMD Ryzen 9950X @ 5.4 RAM: 32GB @ 5800MHz CL26 | BIOS 9902 |
| AMD Alveo X3522 | | 1.12.3.37 |
| DAC | 1 Metre, 10Gbe | |
| Software | | |
| Linux | Rocky Linux release 9.5 (Onyx) | Kernel: 5.14.0-503.31.1.el9_5.x86_64 |
| X3522 Driver | xinlinx_efct | 1.6.6.0 |
| Onload | OpenOnload | 9.0.2.140 |

Tuning

Note: Tuning is identical for “stock” and “overclocked”.

Grub Parameters

```
nomodeset rhgb ipv6.disable=1 rcu_nocb_poll
processor.max_cstate=0 idle=poll nmi_watchdog=0 amd_iommu=off
iomem=relaxed hugepagesz=2MB hugepages=6000 default_hugepagez=2MB
transparent_hugepage=never pcie_aspm=performance audit=0
hpet=disabled nohalt noibrs noibpb nopti edd=off selinux=0
cpuidle.off=1 skew_tick=1 acpi_irq_nobalance nokaslr nospectre_v2
nospectre_v1 l1tf=off nospec_store_bypass_disable no_stf_barrier
mds=off tsx=on tsx_async_abort=off mitigations=off
tsc_calibrate=nocpuid tsc=reliable clocksource=tsc nohz=on
nohz_full=1-15 rcu_nocbs=1-15 nosoftlockup cgroup_disable=memory
skew_tick=1 tsc=reliable rcupdate.rcu_normal_after_boot=1 nohz=on
nohz_full=1-15 rcu_nocbs=1-15 tuned.non_isolcpus=ffff0001
nosoftlockup
```

Tuned

The Tuned service is used to apply AMDULL best practice settings.



More details can be found here: https://www.amd.com/content/dam/amd/en/documents/epyc-technical-docs/white-papers/58649_amd-epyc-tg-low-latency.pdf

Test Parameters

The test script runs 3 tests:

eflatency, sfnt-pingpong [TCP], sfnt-pingpong [UDP]

- Each test is run in sequential byte sizes from 4 bytes to 648 bytes in 4 byte increments.
- Each test is run 3 times.
- Each test allocates an isolated core for the RX and TX part of the test.
- Another 2 isolated cores are assigned to the RX/TX IRQs for the NIC.
- The interfaces on the NIC are isolated from each other using ip namespaces.

Test Code

A snippet of the code is provided below for context of how each test is performed.



```
run_efvi_test () {
    for i in `seq 4 4 648`
    do
        taskset -c ${CORE2} ip netns exec ns1
$ONLOAD_PATH/build/gnu_x86_64/tests/ef_vi/eflatency -s $i pong ${IFACE2} &
        sleep 4
        taskset -c ${CORE1}
$ONLOAD_PATH/build/gnu_x86_64/tests/ef_vi/eflatency -s $i ping ${IFACE1}
    done
}

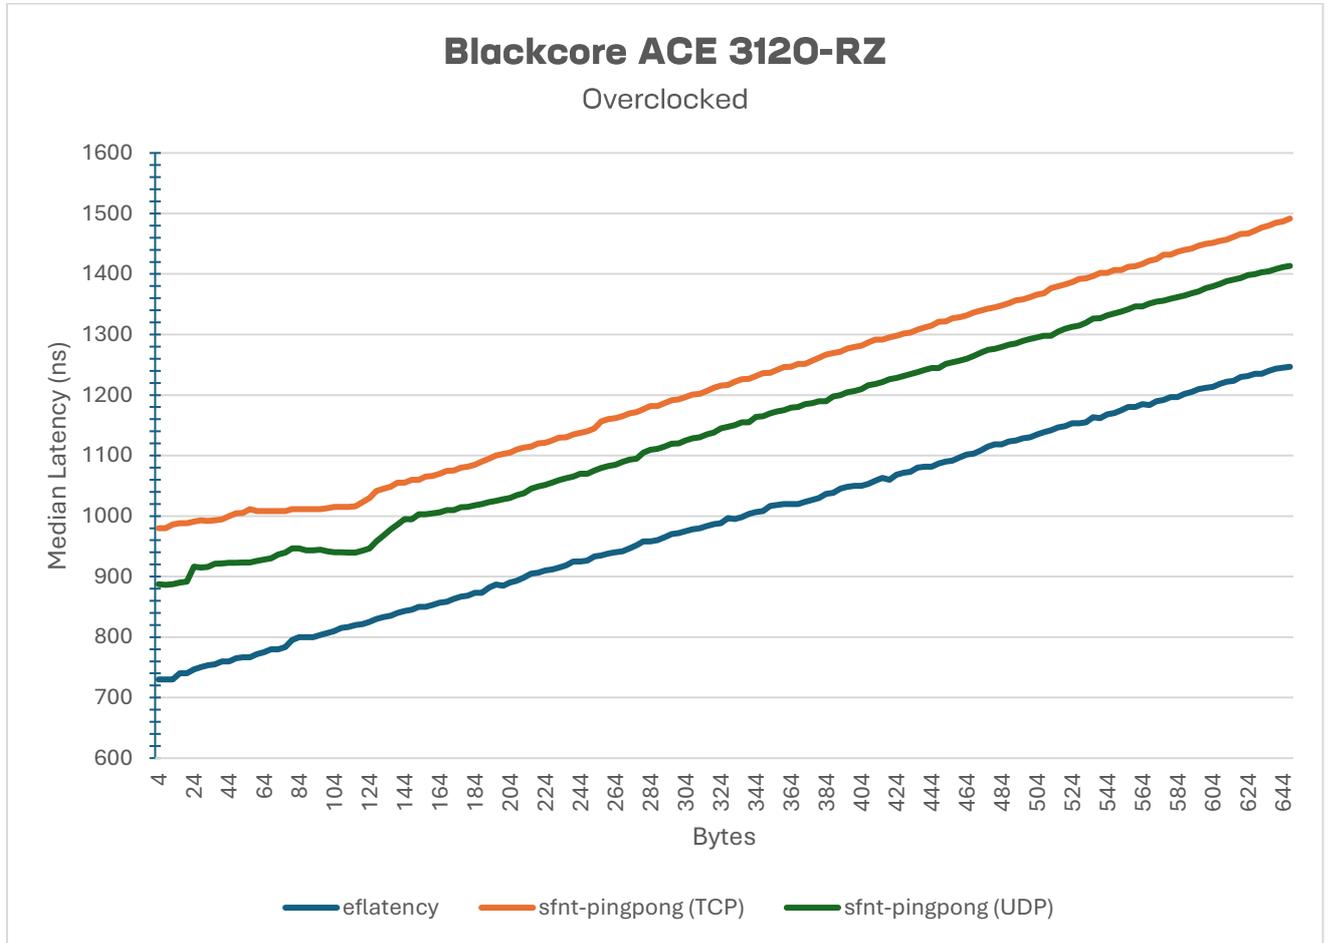
run_tcp_test () {
    SFNTP="/opt/cns-sfnettest/cns-sfnettest-master/src"
    APP="/sfnt-pingpong"
    taskset -c ${CORE2} ip netns exec ns1 onload -p latency-best $SFNTP$APP
&
    taskset -c ${CORE1} onload -p latency-best $SFNTP$APP --affinity
'${CORE1};${CORE1}' --sizes `seq -s , 4 4 648` tcp ${IP2}
}

run_udp_test () {
    SFNTP="/opt/cns-sfnettest/cns-sfnettest-master/src"
    APP="/sfnt-pingpong"
    taskset -c ${CORE2} ip netns exec ns1 onload -p latency-best $SFNTP$APP
&
    taskset -c ${CORE1} onload -p latency-best $SFNTP$APP --affinity
'${CORE1};${CORE1}' --sizes `seq -s , 4 4 648` udp ${IP2}
}
```



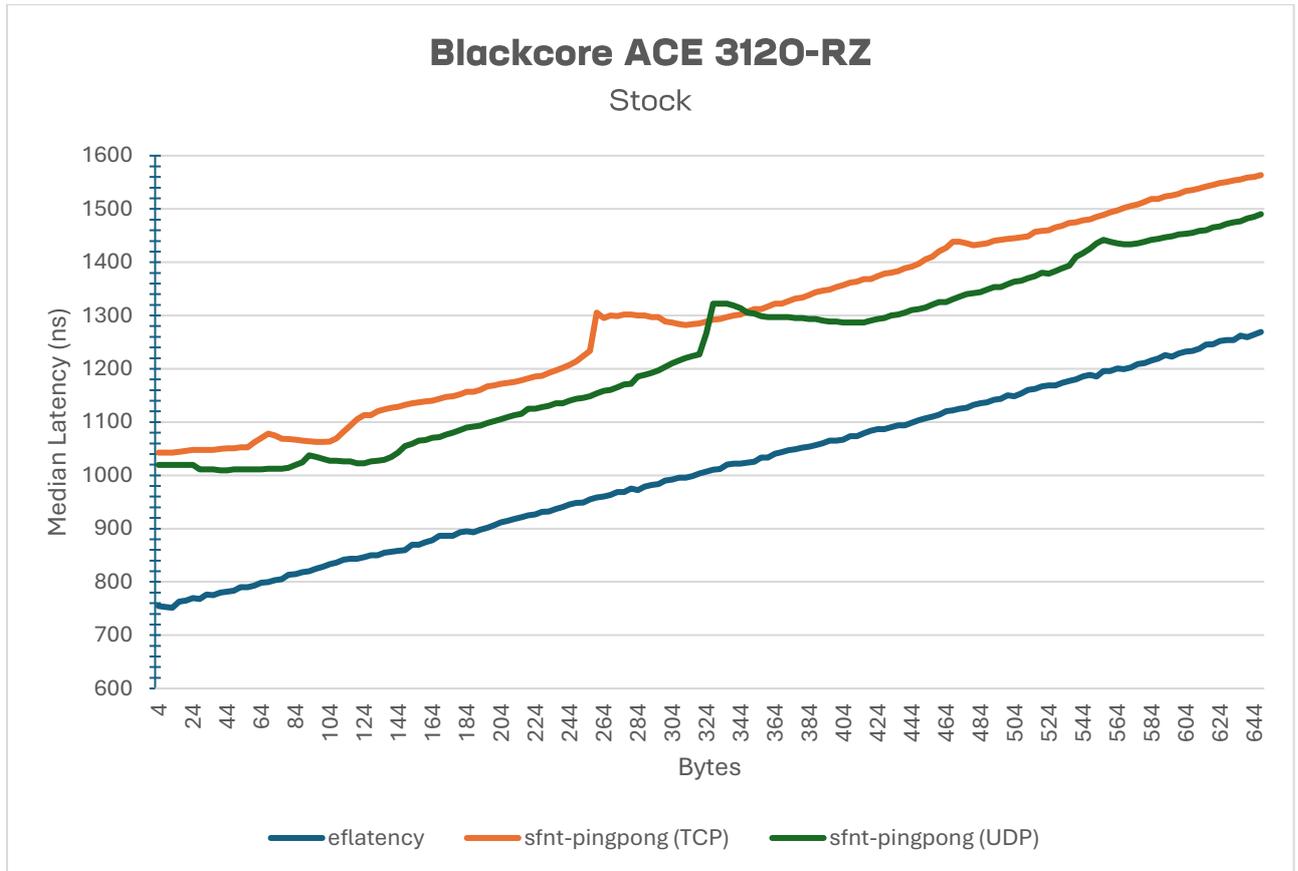
Results

Summary of overclocked results:





Summary of stock results:





Summary comparison at typical byte sizes

4 Bytes

| Min Lat - 4 bytes | | | | |
|----------------------|--|--|--|--|
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |
| | | | | |
| 99%ile Lat - 4 Bytes | | | | |
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |

64 Bytes

Min Lat - 64 bytes

| Test | | | | |
|-----------------------|--|--|--|--|
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |
| | | | | |
| 99%ile Lat - 64 bytes | | | | |
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |



128 Bytes

Min Lat - 128 bytes

| | | | | |
|-------------|--|--|--|--|
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |
| | | | | |

99%ile Lat - 128 bytes

| | | | | |
|-------------|--|--|--|--|
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |

256 Bytes

Min Lat - 256 bytes

| | | | | |
|-------------|--|--|--|--|
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |
| | | | | |

99%ile Lat - 256 bytes

| | | | | |
|-------------|--|--|--|--|
| Test | | | | |
| ef_vi | | | | |
| Onload-TCP | | | | |
| Onload-UDP | | | | |

Conclusion

Blackcore systems provide demonstrable latency improvements over a stock system. Both when looking at minimum achievable latency and significantly more consistent 99%ile performance.



The Blackcore ACE 3120-RZ along with AMD Alveo X3522 measures sub 800ns latency when using ef_vi at 64 byte payloads or less and sub microsecond latency for 256 bytes or less. For UDP latencies with ef_vi, a saving of around 8-13% can be achieved when comparing with stock configuration.

Further information on the **ACE 3120-RZ**:

<https://blackcoretech.com/ace-3100-rz>

Learn more about the Blackcore Labs Program:

See <https://blackcoretech.com/lab-center>

Contact labs@blackcoretech.com

Sales enquiries:

See <https://blackcoretech.com/get-a-quote>

Contact sales@blackcoretech.com