

AOCC compiler (C/C++/Fortran)

Latest release: 3.2, December 2021

<https://developer.amd.com/amd-aocc/>

Architecture	
Generate instructions that runs on AMD 3 rd Gen EPYC™ and AMD 3 rd Gen Ryzen™	-march=znver3
Generate instructions for the local machine	-march=native
Optimization Levels	
Disable all optimizations	-O0
Minimal level speed and code optimization	-O1/ -O
Moderate level optimization	-O2
Aggressive optimizations	-O3
Maximize performance	-Ofast
Enable link time optimization	-flto
Enable loop optimizations	-funroll-loops -enable-licm-vrp -enable-partial-unswitch -fuse-tile-inner-loop -unroll-threshold
Enable advanced loop optimizations	-unroll-aggressive
Enable function level optimizations	-fitodcalls -function-specialize -finline-aggressive -inline-recursion={1..4} (use with flto) -do-block-reordering={none, normal, aggressive}
Enable advanced vectorization	-enable-strided-vectorization -enable-epilog-vectorization
Enable memory layout optimizations	-fremap-arrays (use with -flto)
Profile guided optimizations	-fprofile-instr-generate (1st invocation) -fprofile-instr-use (2nd invocation)
OpenMP®	-fopenmp
For enabling streaming stores, memory bandwidth workloads	-fnt-store
Enable removal of un-used array computation	-reduce-array-computations=3

Other options	
Enable faster, less precise math operations (part of Ofast)	-ffast-math -freciprocal-math
OpenMP® threads and affinity (N number of cores)	export OMP_NUM_THREADS=N export GOMP_CPU_AFFINITY="0-{N-1}"
Enabling vector library	-fveclib=AMDLIBM
Link to AMD library	-L/libm-install-dir/lib -lalm

For Fortran Workloads	
Compile free form FORTRAN	-ffree-form

AMD Optimized Libraries

Latest release: 3.1, Dec 2021

<https://developer.amd.com/amd-aocl/>

AMD µProf (Performance & Power Profiler)

Latest release: 3.4, March 2021

<https://developer.amd.com/amd-uprof/>

AMD EPYC™ 7xx3-series Processors

Compiler Options Quick Reference Guide



GNU compiler collection

Latest release: 11.2, July 2021

Recommended version: GCC trunk, later than 5-Dec-2021

<http://gcc.gnu.org>

Microsoft® Visual Studio 2022

Latest release: 17.0.1.2, November 2021

[https://visualstudio.microsoft.com/
User Guide](https://visualstudio.microsoft.com/User Guide)

Architecture	
Generate instructions that runs on AMD 3 rd Gen EPYC™ and AMD 3 rd Gen Ryzen™	-march=znver3
Generate instructions for the local machine	-march=native
Optimization Levels	
Disable all optimizations (default)	-O0
Minimal level speed and code optimizations	-O1/ -O
Moderate level optimizations	-O2
Aggressive optimizations	-O3
Maximize performance	-Ofast
Additional Optimizations	
Link time optimization	-flto
Enable unrolling	-funroll-all-loops
Generate memory preload instructions	-fprefetch-loop-arrays --param prefetch-latency=300
Profile-guided optimization	-fprofile-generate (1st invocation) -fprofile-use (2nd invocation)
OpenMP®	-fopenmp
Other options	
Enable generation of code that follows IEEE arithmetic	-mieee-fp
Enable faster, less precise math operations	-ffast-math
Compile free form FORTRAN	-ffree-form
OpenMP® threads and affinity (N number of cores)	export OMP_NUM_THREADS=N export GOMP_CPU_AFFINITY="0-{N-1}"
Link to AMD library	-L/libm-install-dir/lib -lalm

Architecture	
Generate instructions that runs on AMD 3 rd Gen EPYC™ and	/arch:[AVX AVX2]
Optimize for 64-bit AMD processors	/favor:AMD64 /d2zveroupper
Optimization Levels	
Disable optimizations	/Od
Maximum optimizations (favor space)	/O1
Maximum optimizations (favor speed)	/O2
[link.exe] Eliminate unreferenced function and/ or data	/OPT:REF
[link.exe] Perform identical COMDAT folding	/OPT:ICF
Output an informational message for loops that are auto-vectorized	/Qvec-report:[1 2]
Enable automatic parallelization of loops, used in conjunction with #pragma loop() directive	/Qpar
Output an informational message for loops that are auto-parallelized	/Qpar-report:[1 2]
Additional Optimizations	
Maintain the precision for floating-point operations through proper rounding	/fp:precise
Optimize floating-point code for speed at the expense of floating-point accuracy and correctness	/fp:fast
Whole Program Optimization (link-time code generation)	/GL
Profile-guided optimization	LTCC:PGI and /LTCC:PGO

Glibc

Latest release: 2.34, August 2021

Recommendation: 2.26 or later

<https://www.gnu.org/software/libc/>

Binutils

Recommendation: 2.26 or later

<https://www.gnu.org/software/binutils/>

Intel® compilers

Latest release: 2021.4

<http://software.intel.com>

Architecture	
Generate instructions that runs on AMD 3 rd Gen EPYC™ and AMD 3 rd Gen Ryzen™	-march=core-avx2 (preferred) OR -axCORE-AVX2
Optimization Levels	
Disable all optimizations	-O0
Speed optimization without code growth	-O1
Enable optimization for speed including vectorization	-O2
Aggressive optimization	-O3
Maximize performance	-Ofast
Additional Optimizations	
Aggressive unrolling	-unroll-aggressive
Disable improved precision floating divides	-no-prec-div
Enable vectorization	-vec
Inter procedural Optimization	-ipo
OpenMP®	-qopenmp
Prefetch optimization	-qopt-prefetch
Profile generated optimization	-prof-gen and -prof-use
Use optimized header definitions	-use-intel-optimized-headers
Other Options	
Floating point accuracy tuning	-fp-model
Compile free form FORTRAN	-free

Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Third-party content is licensed to you directly by the third party that owns the content and is not licensed to you by AMD. ALL LINKED THIRD-PARTY CONTENT IS PROVIDED "AS IS" WITHOUT A WARRANTY OF ANY KIND. USE OF SUCH THIRD-PARTY CONTENT IS DONE AT YOUR SOLE DISCRETION AND UNDER NO CIRCUMSTANCES WILL AMD BE LIABLE TO YOU FOR ANY THIRD-PARTY CONTENT. YOU ASSUME ALL RISK AND ARE SOLELY RESPONSIBLE FOR ANY DAMAGES THAT MAY ARISE FROM YOUR USE OF THIRD-PARTY CONTENT.

ATTRIBUTION

© 2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD EPYC, AMD Ryzen and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. OpenMP, Microsoft, Intel are for informational purposes only and may be trademarks of their respective owners.