AMD EPYC™ 9xx4-series Processors
Compiler Options Quick Reference Guide

AOCC compiler (C/C++/Fortran)

Latest release: 4.0, November 2022

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

Architecture	
Generate instructions that runs on AMD 4th Gen EPYC ™ and AMD 4th Gen Ryzen™	-march=znver4
Generate instructions supported in the given machine	-march=native
Optimization Levels	
Disables all optimizations	-O0
Enables minimal level optimizations	-O1/ -O
Enables moderate level optimizations	-O2
Enables all optimizations that attempt to make programs run faster	-O3
Enables O3 with other aggressive optimizations that may violate strict compliance and precisions	-Ofast
Enables link time optimization	-flto
More advanced optimizations - Enables improved variants of various scalar, vector and loop transformations	-zopt
Enables advanced vector transformations	-fvector-transform -mllvm -enable-strided-vectorization
Enables loop transformations	-floop-transform
Enables advanced loop transformations	-faggressive-loop-transform
Enables memory layout optimizations	-flto -fremap-arrays -mllvm -reduce-array-computations=3
Enables function level optimizations	-flto -fitodcalls -mllvm -function-specialize -flto -finline-aggressive -flto -finline-recursion={1..4}
Profile guided optimizations	-fprofile-instr-generate (1st invocation) -fprofile-instr-use (2nd invocation)
Enables use of OpenMP® directives	-fopenmp
Enables streaming stores to optimize memory bandwidth usage	-fnt-store

Other options	
Enables 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}"
Link to AMD library	-L/libm-install-dir/lib -lamdlibm -lm
Enables vector library	-lamdlibm -fveclib=AMDLIBM -lm
Enables faster library	-lamdlibm -fsclrlib=AMDLIBM -lamdlibmfast -lm
For Fortran Workloads	
Compiles Fortran free form layout	-ffree-form

AMD Optimized Libraries

Latest release: 4.0, November 2022

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

AMD µProf (Performance & Power Profiler)

Latest release: 4.0 November 2022

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

Advanced Micro Devices
One AMD Place, P.O. Box 3453
Sunnyvale, CA 94088-3453

AMD EPYC™ 9xx4-series Processors
Compiler Options Quick Reference Guide

GNU compiler collection

Latest release: GCC 12.2, August 2022
Recommended version: GCC 13 trunk, later than 16th January 2023
http://gcc.gnu.org

Architecture	
Generate instructions that runs on AMD 4th Gen EPYC ™ and AMD 4th Gen Ryzen™	-march=znver4
Generate instructions supported in the given machine	-march=native
Optimization Levels	
Disables all optimizations (default)	-O0
Enables minimal level optimizations	-O1/ -O
Enables moderate level optimizations	-O2
Enables all optimizations that attempt to make programs run faster	-O3
Enables O3 with other aggressive optimizations that may violate strict compliance and precisions	-Ofast
Additional Optimizations	
Enables link time optimizations	-flto
Enables unrolling	-funroll-all-loops
Generates memory preload instructions	-fprefetch-loop-arrays --param prefetch-latency=300
Enables profile-guided optimizations	-fprofile-generate (1st invocation) -fprofile-use (2nd invocation)
Enables use of OpenMP® directives	-fopenmp
Other options	
Enables compiler to use IEEE FP comparisons	-mieee-fp
Enables faster, less precise math operations	-ffast-math
Compiles Fortran free form layout	-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 -lamdlibm -lm

Microsoft® Visual Studio 2022

Latest release: 17.0.15, October 2022
https://visualstudio.microsoft.com/
User Guide

Architecture	
Generate instructions that runs on AMD 4th Gen EPYC ™ and AMD 4th Gen Ryzen™	/arch:[AVX|AVX2]
Optimize for 64-bit AMD processors	/favor:AMD64
Optimization Levels	
Disable optimizations	/Od
Maximum optimizations (favor space)	/O1 [includes /Ob2]
Maximum optimizations (favor speed)	/O2 [includes /Ob2]
Enables inline expansion	/Ob (0/1/2/3)
[link.exe] Eliminates unreferenced function and/ or data	/OPT:REF
[link.exe] Performs identical COMDAT folding	/OPT:ICF
Output an informational message for loops that are auto-vectorized	/Qvec-report:[1|2]
Enables 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
Enables Profile-guided optimizations	LTCG:PGI and /LTCG:PGO

GlibC

Latest release: 2.36, August 2022
Recommendation: 2.26 or later
https://www.gnu.org/software/libc/

Binutils

Latest release: 2.40, January 2023
Recommendation: 2.26 or later
https://www.gnu.org/software/binutils/

Advanced Micro Devices
One AMD Place, P.O. Box 3453
Sunnyvale, CA 94088-3453

Intel® oneAPI DPC++/C++ Compiler

Latest release: 2023.0

http://software.intel.com

Architecture	
Generate instructions that runs on AMD 4th Gen EPYC ™ and AMD 4th Gen Ryzen™	-axCORE-AVX512
Optimization Levels	
Disable all optimizations	-O0
Speed optimization without code growth	-O1
Enables optimization for speed including vectorization	-O2
Enables O2 and aggressive loop	-O3
Enables set of aggressive options to	-Ofast
Additional Optimizations	
Sets function inline level	-inline-level=<value>
Sets unroll loop maximum threshold	-unroll<value>
Disable improved precision floating divides	-no-prec-div
Enables vectorization	-vec
Enables inter procedural optimiza-	-ipo
Enables use of OpenMP® directives	-qopenmp
Enables profile generated optimiza-	-prof-gen and -prof-use
Other Options	
Enables floating point accuracy tun-	-fp-model
Compiles Fortran free form layout	-free

Advanced Micro Devices
One AMD Place, P.O. Box 3453
Sunnyvale, CA 94088-3453

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

Advanced Micro Devices
One AMD Place, P.O. Box 3453
Sunnyvale, CA 94088-3453