

Geekbench 5 CPU Workloads



Introduction	3
Platform Support	4
Compilers	4
Runtime	5
Scores	6
Cryptography Workloads	7
AES-XTS	7
Integer Workloads	8
Text Compression	8
Image Compression	8
Navigation	8
HTML5	8
SQLite	9
PDF Rendering	9
Text Rendering	9
Clang	10
Camera	10
Floating Point Workloads	11
N-Body Physics	11
Rigid Body Physics	11
Gaussian Blur	11
Face Detection	12
Horizon Detection	12
Image Inpainting	12
HDR	13
Ray Tracing	13
Structure from Motion	14
Speech Recognition	14
Machine Learning	14

Introduction

This document outlines the workloads included in the Geekbench 5 CPU Benchmark suite.

CPU Benchmark scores are used to evaluate and optimize CPU and memory performance using workloads that include data compression, image processing, machine learning, and physics simulation. Performance on these workloads is important for a wide variety of applications including web browsers, image editors, and developer tools.

Platform Support

Platform	Minimum Version	Comment
Android	Android 7.0 "Nougat"	
iOS	iOS 12	
Linux	Ubuntu 16.04 LTS	
macOS	macOS 10.13	
Windows	Windows 10	

Compilers

Geekbench 5.0 is built using the following compilers:

Platform	Compiler	Comment
Android	Clang 8.0	Provided by NDK r20
iOS	Xcode 10.3	
Linux	Clang 8.0	
macOS	Xcode 10.3	
Windows	Clang 8.0	

Runtime

Geekbench 5 groups CPU workloads into two sections:

1. Single-Core Workloads
2. Multi-Core Workloads

Each single-core workload has a multi-core counterpart, and vice versa. Each section is grouped into three subsections:

1. Cryptography Workloads
2. Integer Workloads
3. Floating-Point Workloads

Geekbench inserts a pause (or gap) between each workload to minimize the effect thermal issues have on workload performance. Without this gap, workloads that appear later in the benchmark would have lower scores than workloads that appear earlier in the benchmark.

The default gap is 1s for both single-core and multi-core workloads.

Scores

Geekbench 5 scores are calibrated against a baseline score of 1,000 (which is the score of a Dell Precision 3430 with a Core i3-8100 processor). Higher scores are better, with double the score indicating double the performance.

Geekbench 5 provides two composite scores: single-core and multi-core. These scores are computed using a weighted arithmetic mean of the subsection scores. The subsection scores are computed using the geometric mean of the scores of the workloads contained in that subsection.

Subsection	Weight
Cryptography	5%
Integer	65%
Floating Point	30%

Cryptography Workloads

AES-XTS

The Advanced Encryption Standard (AES) defines a symmetric block encryption algorithm. AES encryption is widely used to secure communication channels (e.g., HTTPS) and to secure information (e.g., storage encryption, device encryption).

The AES-XTS workload in Geekbench 5 encrypts a 128MB buffer using AES running in XTS mode with a 256-bit key. The buffer is divided into 4K blocks. For each block, the workload derives an XTS counter using the SHA-1 hash of the block number. The block is then processed in 16-byte chunks using AES-XTS, which involves one AES encryption, two XOR operations, and a $GF(2^{128})$ multiplication.

Geekbench will use AES (including VAES) and SHA-1 instructions when available, and fall back to software implementations otherwise.

Superior AES performance can translate into improved usability for mobile devices. See, e.g., [the Ars Technica review of the Moto E](#).

Integer Workloads

Text Compression

LZMA (Lempel-Ziv-Markov chain algorithm) is a lossless compression algorithm. The algorithm uses a dictionary compression scheme (the dictionary size is variable and can be as large as 4GB). LZMA features a high compression ratio (higher than bzip2).

The LZMA workload compresses and decompresses a 2399KB HTML ebook using the LZMA compression algorithm with a dictionary size of 2048KB. The workload uses the LZMA SDK for the implementation of the core LZMA algorithm.

Image Compression

The Image Compression workload compresses and decompresses a photograph using JPEG, and a CSS sprite using PNG. The workload sets the JPEG quality parameter to “90”, a commonly-used setting for users who desire high-quality images.

The workload uses libjpeg-turbo for the implementation of the core JPEG algorithm, and libpng for the implementation of the core PNG algorithm.

Navigation

The Navigation workload computes driving directions between a sequence of destinations using Dijkstra's algorithm. Similar techniques are used to compute paths in games, to route computer network traffic, and to route driving directions. The dataset contains 216,548 nodes and 450,277 edges with weights approximating travel time along the road represented by the edge. The route includes 13 destinations. The dataset is based on Open Street Map data for Ontario, Canada.

HTML5

The HTML5 workload models DOM creation from both server-side rendered (SSR) and client-side rendered (CSR) HTML5 documents. For the SSR document, the HTML5 workload uses the Gumbo HTML5 parser to create the DOM by parsing an HTML file. For the CSR document, the HTML5 workload uses the Gumbo HTML5 parser to create the DOM by parsing an HTML file, then uses the Duktape JavaScript engine to extend the DOM.

SQLite

SQLite is a self-contained SQL database engine, and is the most widely deployed database engine in the world.

The SQLite workload executes SQL queries against an in-memory database. The database is synthetically created to mimic financial data, and is generated using techniques outlined in “Quickly Generating Billion-Record Synthetic Databases” by J. Gray et al. The workload is designed to stress the underlying engine using a variety of SQL features (such as primary and foreign keys) and query keywords such as: SELECT, COUNT, SUM, WHERE, GROUP BY, JOIN, INSERT, DISTINCT, and ORDER BY. This workload measures the transaction rate a device can sustain with an in-memory SQL database.

PDF Rendering

The Portable Document Format (PDF) is a standard file format used to present and exchange documents independent of software or hardware. PDF files are used in numerous ways, from government documents and forms to e-books.

The PDF workload parses and renders a PDF map of Crater Lake National Park at 200dpi. The PDF workload uses the PDFium library (which is used by Google Chrome to display PDFs).

Text Rendering

The Text Rendering workload parses a Markdown-formatted document and renders it as rich text to a bitmap. The Text Rendering workload uses the following libraries as part of the workload:

- GitHub Flavored Markdown, used to parse the Markdown document.
- FreeType, used to render fonts.
- ICU (International Components for Unicode), used for boundary analysis.

The Text Rendering workload input file is 1721 words long, and produces a bitmap that is 1275 pixels by 9878 pixels in size.

Clang

Clang is a compiler front end for the programming languages C, C++, Objective-C, Objective-C++, OpenMP, OpenCL, and CUDA. It uses LLVM as its back end.

The Clang workload compiles a 1,094 line C source file (of which 729 lines are code). The workload uses AArch64 as the target architecture for code generation.

Camera

Camera replicates a photo sharing application like Instagram. Camera merges several steps into one workload:

- SHA2 checksum generation
- JSON parsing
- Image compositing
- Image filters (gaussian blur, contrast)
- Image resizing (thumbnail generation)
- SQLite (SELECT images to be processed)

All steps run on the CPU and are not accelerated by the GPU.

Floating Point Workloads

N-Body Physics

The N-Body Physics workload computes a 3D gravitation simulation using the Barnes-Hut method. To compute the exact gravitational force acting on a particular body x in a field of N bodies requires $N - 1$ force computations. The Barnes-Hut method reduces the number of force computations by approximating as a single body any tight cluster of bodies that is far away from x . It does this efficiently by dividing the space into octants — eight cubes of equal size — and recursively subdividing each octant into octants, forming a tree, until each leaf octant contains exactly one body. This recursive subdivision of the space requires floating point operations and non-contiguous memory accesses.

The N-Body Physics workload operates on 16,384 bodies arranged in a “flat” galaxy with a massive black hole in its centre.

Rigid Body Physics

The Rigid Body Physics workload computes a 2D physics simulation for rigid bodies that includes collisions and friction. The workload uses the Lua programming language to initialize and manage the physics simulation, and uses the Box2D physics library to perform the actual physics calculations.

Gaussian Blur

The Gaussian Blur workload blurs an image using a Gaussian spatial filter. Gaussian blurs are widely used in software, both in operating systems to provide interface effects, and in image editing software to reduce detail and noise in an image. Gaussian blurs are also used in computer vision applications to enhance image structures at different scales.

The Gaussian Blur workload blurs a 24 megapixel image using a Gaussian spatial filter. While the Gaussian blur implementation supports an arbitrary sigma, the workload uses a fixed sigma of 3.0f. This sigma translates into a filter diameter of 25 pixels by 25 pixels.

Face Detection

Face detection is a computer vision technique that identifies human faces in digital images. One application of face detection is in photography, where camera applications use face detection for autofocus.

The Face Detection workload uses the algorithm presented in “Rapid Object Detection using a Boosted Cascade of Simple Features” (2001) by Viola and Jones. The algorithm can produce multiple boxes for each face. These boxes are reduced to a single box using non-maximum suppression.

Horizon Detection

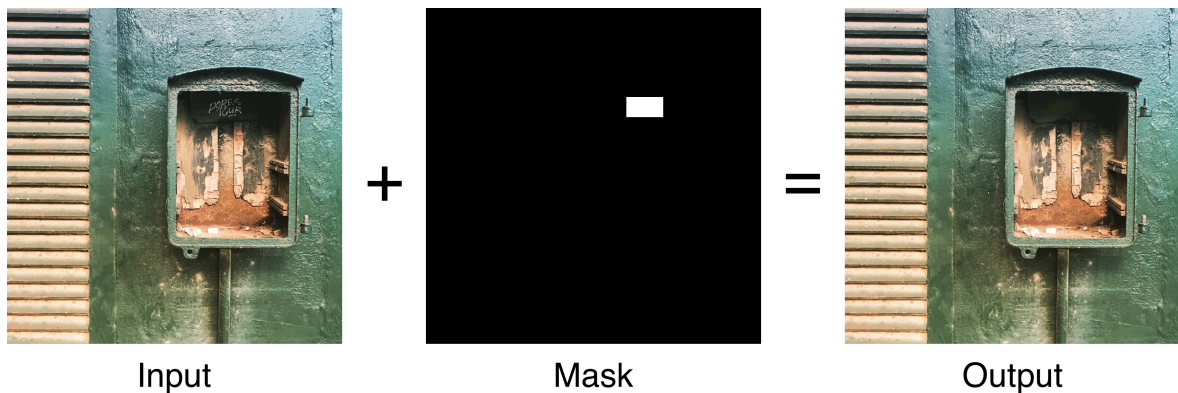
The Horizon Detection workload searches for the horizon line in an image. If the horizon line is found, the workload rotates the image to make the horizon line level.

The workload first applies a Canny edge detector to the image to reduce details, then detects lines in the image using the Hough transform, and then picks the line with the maximum score as the horizon. The workload rotates the image so the horizon line is level in the image.

Image Inpainting

The Image Inpainting workload takes an input image with an undesirable region (indicated via a mask image) and uses an inpainting scheme to reconstruct the region using data from outside the undesirable region.

The Image Inpainting workload operates on 1 megapixel images.



HDR

The HDR workload takes four standard dynamic range (SDR) images and produces a high dynamic range (HDR) image.

The HDR workload uses the algorithm described in the paper, "Dynamic Range Reduction inspired by Photoreceptor Physiology" by Reinhard and Devlin, and produces superior images as compared to the tone mapping algorithm in Geekbench 4.

Ray Tracing

Ray tracing is a rendering technique. Ray tracing generates an image by tracing the path of light through an image plane and simulating the effects of its encounters with virtual objects. This method is capable of producing high-quality images, but these images come at a high computational cost.

The Ray Tracing workload uses a k-d tree, a space-partitioning data structure, to accelerate the ray intersection calculations.

The Ray Tracing workload operates on a scene with 3,608 textured triangles. The rendered image is 768 pixels by 768 pixels.



Structure from Motion

Augmented reality (AR) systems add computer-generated graphics to real-world scenes. The systems must have an understanding of the geometry of the real-world scene in order to properly integrate the computer-generated graphics. One approach to calculating the geometry is through Structure from Motion (SfM) algorithms.

The Structure from Motion workload takes two 2D images of the same scene and constructs an estimate of the 3D coordinates of the points visible in both images.

Speech Recognition

The Speech Recognition workload performs recognition of arbitrary English speech using PocketSphinx, a widely used library that uses HMM (Hidden Markov Models).

Using speech to interact with smartphones is becoming more popular with the introduction of Siri, Google Assistant, and Cortana, and this workload tests how quickly a device can process sound and recognize the words that are being spoken.

Machine Learning

The Machine Learning workload is an inference workload that executes a Convolutional Neural Network to perform an image classification task. The workload uses MobileNet v1 with an alpha of 1.0 and an input image size of 224 pixels by 224 pixels. The model was trained on the ImageNet dataset.