Ryzen AI

Release 1.8.0

AMD

Oct 07, 2026

CONTENTS

1 Quick Start 3

2 Development Flow Overview 5

3 LLM Flow Overview 7

AMD Ryzen™ AI Software includes the tools and runtime libraries for optimizing and deploying AI inference on AMD Ryzen™ AI powered PCs. Ryzen AI software enables applications to run on the neural processing unit (NPU) built in the AMD XDNA™ architecture, as well as on the integrated GPU. This allows developers to build and deploy models trained in PyTorch or TensorFlow and run them directly on laptops powered by Ryzen AI using ONNX Runtime and the Vitis™ AI Execution Provider (EP).

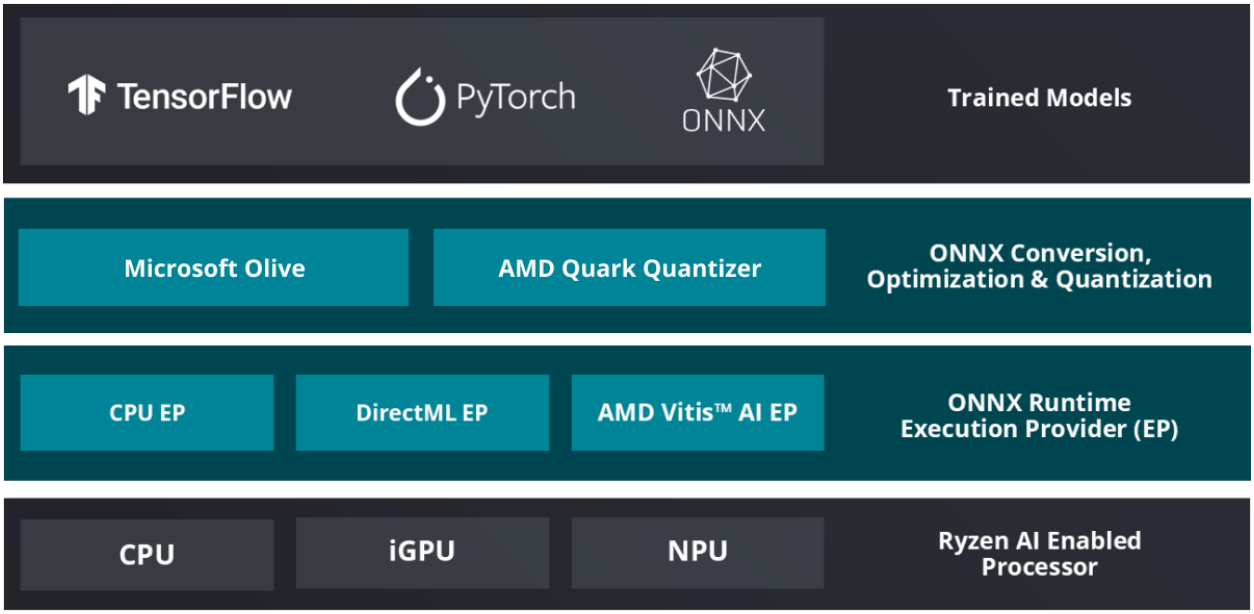

CHAPTER

ONE

QUICK START

- *Supported Configurations*
- *Installation Instructions*
- *Examples, Demos, Tutorials*

CHAPTER

TWO

DEVELOPMENT FLOW OVERVIEW

The Ryzen AI development flow does not require any modifications to the existing model training processes and methods. The pre-trained model can be used as the starting point of the Ryzen AI flow.

2.1 Quantization

Quantization involves converting the AI model's parameters from floating-point to lower-precision representations, such as 8-bit integer. Quantized models are more power-efficient, utilize less memory, and offer better performance. Ryzen AI Software also supports CNN and Transformer models in floating-point 32 format as input models without quantization. These models are internally converted to bfloat16 and compiled using the bfloat16 compilation flow.

AMD Quark is a comprehensive cross-platform deep learning toolkit designed to simplify and enhance the quantization of deep learning models. Supporting both PyTorch and ONNX models, Quark empowers developers to optimize their models for deployment on a wide range of hardware backends, achieving significant performance gains without compromising accuracy.

For more details, refer to the *Model Quantization* page.

2.2 CNN/Transformer Compilation and Deployment

The AI model is deployed using the ONNX Runtime with either C++ or Python APIs. The Vitis AI Execution Provider included in the ONNX Runtime intelligently determines what portions of the AI model should run on the NPU, optimizing workloads to ensure optimal performance with lower power consumption.

For more details, refer to the *Model Compilation and Deployment* page.

CHAPTER

THREE

LLM FLOW OVERVIEW

The Ryzen AI LLM software stack is available through three development interfaces, each suited for specific use cases as outlined in the sections below. All three interfaces are built on top of native OnnxRuntime GenAI (OGA) libraries or llama.cpp libraries, as shown in the *Ryzen AI Software Stack* diagram below.

The **high-level Python APIs**, as well as the **Server Interface**, also leverage the **Lemonade SDK**, which is multi-vendor open-source software that provides everything necessary for quickly getting started with LLMs on OGA or llama.cpp.

At the bottom, **OnnxRuntime GenAI (OGA)** or llama.cpp (only supported for iGPU) API is the lowest-level API available for building LLM applications on a Ryzen AI PC.

Table 1: Ryzen AI Software Stack

Your Python Application	Your LLM Stack	Your Native Application
*Lemonade Python API**	*Lemonade Server Interface**	OGA C++ Headers **OR** llama.cpp C++ Headers
Custom AMD OnnxRuntime GenAI (OGA) **OR** llama.cpp*		
AMD Ryzen AI Driver and Hardware		

For more details, refer to the *LLM Deployment Overview* page.

3.1 Release Notes

3.1.1 Supported Configurations

Ryzen AI 1.8 Software supports AMD processors codenamed Phoenix, Hawk Point, Strix, Strix Halo, and Krackan Point. These processors can be found in the following Ryzen series:

- Ryzen 200 Series
- Ryzen 7000 Series, Ryzen PRO 7000 Series
- Ryzen 8000 Series, Ryzen PRO 8000 Series
- Ryzen AI 300 Series, Ryzen AI PRO Series, Ryzen AI Max 300 Series

For a complete list of supported devices, refer to the processor specifications page (look for the "AMD Ryzen AI" column towards the right side of the table, and select "Available" from the pull-down menu).

The rest of this document will refer to Phoenix as PHX, Hawk Point as HPT, Strix and Strix Halo as STX, and Krackan Point as KRK.

3.1.2 Model Compatibility Table

The following table lists which types of models are supported on what hardware platforms.

Model Type	PHX/HPT	STX/KRK
CNN INT8		
CNN BF16		
NLP BF16		
LLM (OGA)		

3.1.3 Version 1.8

- LLM
  - 8 new LLM and VLM support: Hunyuan-0.5B-Instruct, Hunyuan-1.8B-Instruct, Hunyuan-4B-Instruct, Hunyuan-7B-Instruct, HY-MT1.5-1.8B, HY-MT1.5-7B, VideoChat-Flash-Qwen2_5-7B_InternVideo2-1B, Qwen3-VL-4B-Instruct
  - New TTS model support: neutts-air, neutts-nano
  - New embedding model support: embeddinggemma-300m
  - New translation model support: translategemma-4b-it
- Stable Diffusion
  - New models, new applications, more dynamic resolutions

        * SD1.5 i2i-canny
        * SDXL-turbo t2i with 5x Dynamic Resolutions
        * SDXL-base t2i with 20x Dynamic Resolutions
        * SSD-1B t2i with 20x Dynamic Resolutions
        * Playground v2.5 t2i with 20x Dynamic Resolutions
        * DreamShaper-XL-Lightning t2i with 20x Dynamic Resolutions
        * FLUX.1-Schnell t2i with 20x Dynamic Resolutions
        * SD3.0-medium t2i with 20x Dynamic Resolutions
        * SD3.5-medium t2i with 20x Dynamic Resolutions
        * SD3.0-inpainting
        * Flux.2-Klein-4B t2i
    - Performance optimizations with -O1 for SD3.x models
    - Support running large models on low-memory machine
- Whisper.cpp: Improved short-audio transcription speed, with 3-second audio RTF reduced to 13.9% (Whisper small) and 33.7% (Whisper large-v3-turbo)
- Known Issues
    - LLM
        * The following models are not supported in the current release:
            · Full Fusion: Qwen2.5-14B, Qwen2.5-14B-Instruct, Qwen3-0.6B, Qwen3-14B
            · TPS Fusion: Phi-3-mini-128k-instruct is only supported on Windows
        * Model generation is not supported on Linux in this release. Models generated on Windows are compatible with Linux
    - BF16 Models
        * The `openai-whisper-medium-encoder` model requires the following additional compiler option `FE_EXPERIMENT+=enable-binary-blob-inputs=0` in the BF16 config file

3.1.4 Version 1.7.1

- LLM
    - Performance improvements for Gemma-3-4B and GPT-OSS-20B
    - NPU support for SmolLM2-135M-Instruct and SmolLM-135M-Instruct
    - Long context support up to 16K for NPU models

- Stable Diffusion
  - Support for 10 model classes (t2i, i2i, dynamic shapes, ControlNet)
  - Performance and memory improvements for SD-2.1-v and SD-3.x
- Ryzen AI installer is also available on Linux platform.

3.1.5 Version 1.7

- CNN/Transformer
  - General Performance improvements (Avg 15-20% improvements expected)
  - ~40% Compile time improvement for transformer models
  - Whisper.cpp integration enabling scalable ASR support
- LLM
  - New LLM Support: Qwen-2.5-14b-Instruct, Qwen-3-14b-Instruct, Phi-4-mini-instruct
  - First preview release of Sparse-LLM: GPT-OSS-20b NPU model
  - First preview release of VLM: Gemma-3-4b-it NPU model (functional)
  - Long context support for hybrid execution models
- Stable Diffusion
  - New model support: SD3.5-Turbo with 8x dynamic resolutions and 2x dynamic batches (Text2Image and Image2ImageControlNet)
  - New model support: Segmind-Vega 1024x1024 (Text2Image)
  - Up to 40% performance improvement for all supported models (Native-BFP16 format)
- Breaking Changes
  - The `log_level` provider option is deprecated. Use the ONNX Runtime session option `log_severity_level` instead to control Vitis AI execution logs.

3.1.6 Version 1.6

- BF16 Compiler (CNN, Transformer, ASR)
  - BF16 CNN perf improvements average 80% across release
  - BF16 perf improvements - 1.3X faster on CNN than iGPU and 2.6X faster for transformers than iGPU
  - Improved coverage and improved performance for ASR models
  - Average 3x compile time improvement
  - Smaller installation size

  - Reduction in CPU overhead by pushing data layout transformation to NPU
  - Dynamic batch size support for compilation
- New Integer Compiler (CNN)
  - Support for General Asymmetric Quantization enabling third party quantized models to run on NPU
  - Support for XINT8, A8W8, A16W8
- LLM
  - Broad Set of NPU only models optimized performance
  - New set of hybrid models with bfp16 activation
  - New architecture support in hybrid flow (Phi-4, Qwen-3)
  - Context length improvement from 2K to 4K for all models.
- Stable Diffusion Demo
  - 8x Dynamic Resolution for SD3.0/3.5 (text2image and image2imageControlNet)
  - Performance boost for SD 1.5/2.1-base/turbo/XL-turbo
  - Support Batch Size 1 for SD-turbo/SDXL-turbo
  - New model support (SD2.1-v 768x768 text2image, SDXL-base 1024x1024 text2image)
- Breaking Changes
  - For running INT8 models on STX/KRK or newer devices, the `xclbin` provider option is no longer supported and should no longer be used. See *Using INT8 Models* for full details.
  - For running INT8 models on PHX/HPT devices, the `target` option should be set to `X1`. The NPU binary should still be specified using the `xclbin` provider option. See *Using INT8 Models* for full details.
  - For BF16 models, the default configurations file requires a new `target` section. See *Config File Options* for full details.
  - LLM:
    * OGA version has been updated to v0.9.2 (Ryzen AI 1.6) from v0.7.0 (Ryzen AI 1.5). Any APIs that are obsolete must be updated to the supported equivalents as described in the Microsoft ONNX Runtime GenAI v0.9.2 documentation
- Hybrid models published with earlier releases are not compatible with Ryzen AI 1.6. use the hybrid models published with the 1.6 release.

3.1.7 Version 1.5

- EoU Improvement
  - Application concurrency: improves the resource distribution across applications
  - Model Compilation time: 2x – 8x faster
  - Installation Size: 80% smaller
- Stable Diffusion demo pipelines (preview)
- 4K context length supported (on selected models)
- LLM Context cache support (on selected models)
- Bug fixes
- New LLMs released
  - Qwen/Qwen2.5-1.5B-Instruct
  - Qwen/Qwen2.5-3B-Instruct
  - Qwen/Qwen2.5-7B-Instruct
- Breaking Changes
  - The `%RYZEN_AI_INSTALLATION_PATH%\deployment` folder has been reorganized and flattened. Deployment DLLs are no longer organized in subfolders. If you use application build scripts that pull DLLs from the `deployment` folder, you need to update them based on the new paths. Refer to the *Application Packaging Requirements* section for further details.
  - The `1x4.xclbin` (PHX/HPT) and `AMD_AIE2P_Nx4_Overlay.xclbin` (STX/KRK) NPU binaries are no longer supported and should not be used. You should use the `4x4.xclbin` (PHX/HPT) and `AMD_AIE2P_4x4_Overlay.xclbin` (STX/KRK) NPU binaries instead.
  - The `XLNX_ENABLE_CACHE`, `XLNX_VART_FIRMWARE`, and `XLNX_TARGET_NAME` environment variables are no longer supported and should not be relied upon.
  - Support for VitisAI EP cache encryption is no longer available. To encrypt the compiled models, use the ONNX Runtime *EP Context Cache* feature instead.
  - For INT8 models, the VitisAI EP does not save the compiled model to disk by default. To save the compiled model, use the ONNX Runtime *EP Context Cache* feature or set the `enable_cache_file_io_in_mem` provider option to 0.
  - Generation of the `vitisai_ep_report.json` file is no longer automatic and should be manually enabled. See the *Operator Assignment Report* section for details.
  - Changes to the OGA flow for LLMs:
    - OGA Version is updated to **v0.7.0** (Ryzen AI 1.5) from v0.6.0 (Ryzen AI 1.4).

* The `hybrid_llm` and `npu_llm` folders are consolidated into a new folder named `LLM`, which contains the `model_benchmark.exe` and `run_model.py` scripts, along with the necessary C++ headers and .lib files to support both the Hybrid LLM and NPU LLM workflows in C++ and Python.
* For NPU LLM models, the `vaip_llm.json` file is no longer required. As a result, the `vaip_llm.json` path is removed from the `genai_config.json` for all NPU models. Ensure that you re-download the NPU models from Hugging Face when using the Ryzen AI 1.5 installer.

3.1.8 Version 1.4

- New Features:
  - New architecture support for Ryzen AI 300 series processors
  - Unified support for LLMs, INT8, and BF16 models in a single release package
  - Public release for compilation of BF16 CNN and NLP models on Windows
  - Public release of the LLM Hybrid OGA flow
  - LLM building flow for finetuned LLM
  - Support for up to 16 hardware contexts on Ryzen AI 300 series processors
  - Vitis AI EP now supports the ONNX Runtime EP context cache feature (for custom handling of pre-compiled models)
  - Ryzen AI environment variables converted to VitisAI EP session options
  - Improved exception handling and fallback to CPU
- New Hybrid execution mode LLMs:
  - DeepSeek-R1-Distill-Llama-8B
  - DeepSeek-R1-Distill-Qwen-1.5B
  - DeepSeek-R1-Distill-Qwen-7B
  - Gemma2-2B
  - Qwen2-1.5B
  - Qwen2-7B
  - AMD-OLMO-1B-SFT-DPO
  - Mistral-7B-Instruct-v0.1
  - Mistral-7B-Instruct-v0.2
  - Mistral-7B-v0.3
  - Llama3.1-8B-Instruct

  - Codellama-7B-Instruct

- *New BF16 model examples*:
  - Image classification
  - Finetuned DistilBERT for text classification
  - Text embedding model Alibaba-NLP/gte-large-en-v1.5

- New Tools:
  - Lemonade SDK
    * Lemonade Server: A server interface that uses the standard Open AI API, allowing applications in any language to integrate with Lemonade Server for local LLM deployment and compatibility with existing Open AI apps.
    * Lemonade Python API: Offers High-Level API for easy integration of Lemonade LLMs into Python applications and Low-Level API for custom experiments with specific checkpoints, devices, and tools.
    * Lemonade Command Line Interface easily benchmark, measure accuracy, prompt or gather memory usage of your LLM.
  - TurnkeyML – Open-source tool that includes low-code APIs for general ONNX workflows.
  - Digest AI – A Model Ingestion and Analysis Tool in collaboration with the Linux Foundation.
  - GAIA – An open-source application designed for the quick setup and execution of generative AI applications on local PC hardware.

- Quark-torch:
  - Added ROUGE and METEOR evaluation metrics for LLMs
  - Support for evaluating ONNX models exported using OGA
  - Support for offline evaluation (evaluation without generation) for LLMs
  - Support for Hugging Face integration
  - Support for Gemma2 quantization using the OGA flow
  - Support for Llama-3.2 quantization with FP8 (weights, activation, and KV-cache) for the vision and language components

- Quark-onnx:
  - Support compatibility with ONNX Runtime version 1.20.0, and 1.20.1
  - Support for microexponents (MX) data types, including MX4, MX6, and MX9
  - Support for BF16 data type for VAIML

  - Support for excluding pre and post-processing from quantization
  - Support for mixed precision with any data type
  - Support for Quarot rotation R1 algorithm
  - Support for microexponents and microscaling AdaQuant
  - Support for an auto-search algorithm to automatically find the best accuracy quantized model
  - Added tools for evaluating L2, PSNR, VMAF, and cosine
- ONNX Runtime EP:
  - Support for Chinese characters in the `filename/cache_dir/cache_key/xclbin`
  - Support for `int4/uint4` data type
  - Support for configurable failure handling: CPU fallback or exception
  - Update for encrypt/decrypt feature
- Known Issues:
- Microsoft Windows Insider Program (WIP) users might see warnings or need to restart when running all applications concurrently.
  - NPU driver and workloads will continue to work.
- Context creation might appear to be limited when some application do not close contexts quickly.

3.1.9 Version 1.3

- New Features:
  - Initial release of the Quark quantizer
  - Support for mixed precision data types
  - Compatibility with Copilot+ applications
- Improved support for *LLMs using OGA*
- New EoU Tools:
  - CNN profiling tool for VAI-ML flow
  - Idle detection and suspension of contexts
  - Rebalance feature for AIE hardware resource optimization
- NPU and Compiler:
  - New Op Support:

        * MAC
        * QResize Bilinear
        * LUT Q-Power
        * Expand
        * Q-Hsoftmax
        * A16 Q-Pad
        * Q-Reduce-Mean along H/W dimension
        * A16 Q-Global-AvgPool
        * A16 Padding with non-zero values
        * A16 Q-Sqrt
        * Support for XINT8/XINT16 MatMul and A16W16/A8W8 Q-MatMul
    - Performance Improvements:
        * Q-Conv, Q-Pool, Q-Add, Q-Mul, Q-InstanceNorm
        * Enhanced QDQ support for a range of operations
        * Enhanced the tiling algorithm
        * Improved graph-level optimization with extra transpose removal
        * Enhanced AT/MT fusion
        * Optimized memory usage and compile time
        * Improved compilation messages
- Quark for PyTorch:
    - Model Support:
        * Examples of LLM PTQ, such as Llama3.2 and Llama3.2-Vision models
        * Example of YOLO-NAS detection model PTQ/QAT
        * Example of SDXL v1.0 with weight INT8 activation INT8
    - PyTorch Quantizer Enhancements:
        * Partial model quantization by user configuration under FX mode
        * Quantization of ConvTranspose2d in Eager Mode and FX mode
        * Advanced Quantization Algorithms with auto-generated configurations
        * Optimized Configuration with DataTypeSpec for ease of use
        * Accelerated in-place replacement under Eager Mode

        * Loading configuration from file of algorithms and pre-optimizations

- Quark for ONNX:
  - New Features:
    * Compatibility with ONNX Runtime version 1.18, 1.19
    * Support for int4, uint4, Microscaling data types
    * Quantization for arbitrary specified operators
    * Quantization type alignment of element-wise operators for mixed precision
    * ONNX graph cleaning
    * Int32 bias quantization
  - ONNX Quantizer Enhancements:
    * Fast fine-tuning support for the MatMul operator, BFP data type, and GPU acceleration
    * Improved ONNX quantization of LLM models
    * Optimized quantization of FP16 models
    * Custom operator compilation process
    * Default parameters for auto mixed precision
    * Optimized Ryzen AI workflow by aligning with hardware constraints of the NPU
- ONNX Runtime EP:
  - Support for ONNX Runtime EP shared libraries
  - Python dependency removal
  - Memory optimization during the compile phase
  - Pattern API enhancement with multiple outputs and commutable arguments support
- Known Issues:
  - Extended compile time for some models with BF16/BFP16 data types
- LLM models with a 4K sequence length might revert to CPU execution. - Accuracy drop in some Transformer models using BF16/BFP16 data types, requiring Quark intervention

3.1.10 Version 1.2

- New features:
  - Support added for Strix Point NPUs
  - Support added for integrated GPU

  - Smart installer for Ryzen AI 1.2
  - NPU DPM based on power slider
- New model support:
  - LLM flow support for multiple models in both PyTorch and ONNX flow (optimized model support will be released asynchronously)
  - SDXL-T with limited performance optimization
- New EoU tools:
  - AI Analyzer : Analysis and visualization of model compilation and inference profiling
  - Platform/NPU inspection and management tool (xrt-smi)
  - Onnx Benchmarking tool
- New Demos:
  - NPU-GPU multi-model pipeline application demo
- NPU and Compiler
  - New device support: Strix Nx4 and 4x4 Overlay
  - New Op support:
    * InstanceNorm
    * Silu
    * Floating scale quantization operators (INT8, INT16)
  - Support new rounding mode (Round to even)
  - Performance Improvement:
    * Reduced the model compilation time
    * Improved instruction loading
    * Improved synchronization in large overlay
    * Enhanced strided_slice performance
    * Enhanced convolution MT fusion
    * Enhanced convolution AT fusion
    * Enhanced data movement op performance
- ONNX Quantizer updates
  - Improved usability with various features and tools, including weights-only quantization, graph optimization, dynamic shape fixing, and format transformations.

  - Improved the accuracy of quantized models through automatic mixed precision and enhanced AdaRound and AdaQuant techniques.
  - Enhanced support for the BFP data type, including more attributes and shape inference capability.
  - Optimized the NPU workflow by aligning with the hardware constraints of the NPU.
  - Supported compilation for Windows and Linux.
  - Bugfix:
    * Fixed the problem where per-channel quantization is not compatible with onnxruntime 1.17.
    * Fixed the bug of CLE when conv with groups.
    * Fixed the bug of bias correction.
- Pytorch Quantizer updates
  - Tiny value quantization protection.
  - Higher onnx version support in quantized model exporting.
  - Relu6 hardware constrains support.
  - Support of mean operation with keepdim=True.
- Resolved issues:
  - NPU SW stack will fail to initialize when the system is out of memory. This could impact camera functionality when Microsoft Effect Pack is enabled.
  - If Microsoft Effects Pack is overloaded with other 4+ applications that use NPU to do inference, then camera functionality can be impacted. Can be fixed with a reboot. This will be fixed in the next release.

3.1.11 Version 1.1

- New model support:
  - Llama 2 7B with w4abf16 (3-bit and 4-bit) quantization (Beta)
  - Whisper base (EA access)
- New EoU tools:
  - CNN Benchmarking tool on RyzenAI-SW Repo
  - Platform/NPU inspection and management tool

Quantizer

- ONNX Quantizer:
  - Improved usability with various features and tools, including diverse parameter configurations, graph optimization, shape fixing, and format transformations.
  - Improved quantization accuracy through the implementation of experimental algorithmic improvements, including AdaRound and AdaQuant.
  - Optimized the NPU workflow by distinguishing between different targets and aligning with the hardware constraints of the NPU.
  - Introduced new utilities for model conversion.
- PyTorch Quantizer:
  - Mixed data type quantization enhancement and bug fix.
  - Corner bug fixes for add, sub, and conv1d operations.
  - Tool for converting the S8S8 model to the U8S8 model.
  - Tool for converting the customized Q/DQ to onnxruntime contributed Q/DQ with the "microsoft" domain.
  - Tool for fixing a dynamic shapes model to fixed shape model.
- Bug fixes
  - Fix for incorrect logging when simulating the LeakyRelu alpha value.
  - Fix for useless initializers not being cleaned up during optimization.
  - Fix for external data cannot be found when using use_external_data_format.
  - Fix for custom Ops cannot be registered due to GLIBC version mismatch

NPU and Compiler

- New op support:
  - Support Channel-wie Prelu.
  - Gstiling with reverse = false.
- Fixed issues:
  - Fixed Transpose-convolution and concat optimization issues.
  - Fixed Conv stride 3 corner case hang issue.
- Performance improvement:
  - Updated Conv 1x1 stride 2x2 optimization.
  - Enhanced Conv 7x7 performance.

  - Improved padding performance.
  - Enhanced convolution MT fusion.
  - Improved the performance for NCHW layout model.
  - Enhanced the performance for eltwise-like op.
  - Enhanced Conv and eltwise AT fusion.
  - Improved the output convolution/transpose-convolution's performance.
  - Enhanced the logging message for EoU.

ONNX Runtime EP

- End-2-End Application support on NPU
  - Enhanced existing support: Provided high-level APIs to enable seamless incorporation of pre/post-processing operations into the model to run on NPU
  - Two examples (resnet50 and yolov8) published to demonstrate the usage of these APIs to run end-to-end models on the NPU
- Bug fixes for ONNXRT EP to support customers' models

Misc

- Contains mitigation for the following CVEs: CVE-2024-21974, CVE-2024-21975, CVE-2024-21976

3.1.12 Version 1.0.1

- Minor fix for Single click installation without given env name.
- Perform improvement in the NPU driver.
- Bug fix in elementwise subtraction in the compiler.
- Runtime stability fixes for minor corner cases.
- Quantizer update to resolve performance drop with default settings.

3.1.13 Version 1.0

Quantizer

- ONNX Quantizer
  - Support for ONNXRuntime 1.16.

- Support for the Cross-Layer-Equalization (CLE) algorithm in quantization, which can balance the weights of consecutive Conv nodes to make it more quantize-friendly in per-tensor quantization.
- Support for mixed precision quantization including UINT16/INT16/UINT32/INT32/FLOAT16/BFLOAT16, and support asymmetric quantization for BFLOAT16.
- Support for the MinMSE method for INT16/UINT16/INT32/UINT32 quantization.
- Support for quantization using the INT16 scale.
- Support for unsigned ReLU in symmetric activation configuration.
- Support for converting Float16 to Float32 during quantization.
- Support for converting NCHW model to NHWC model during quantization.
- Support for two more modes for MinMSE for better accuracy. The "All" mode computes the scales with all batches while the "MostCommon" mode computes the scale for each batch and uses the most common scales.
- Support for the quantization of more operations:
  * PReLU, Sub, Max, DepthToSpace, SpaceToDepth, Slice, InstanceNormalization, and LpNormalization.
  * Non-4D ReduceMean.
  * Leakyrelu with arbitrary alpha.
  * Split by converting it to Slice.
- Support for op fusing of InstanceNormalization and L2Normalization in NPU workflow.
- Support for converting Clip to ReLU when the minimal value is 0.
- Updated shift_bias, shift_read, and shift_write constraints in the NPU workflow and added an option "IPULimitationCheck" to disable it.
- Support for disabling the op fusing of Conv + LeakyReLU/PReLU in the NPU workflow.
- Support for logging for quantization configurations and summary information.
- Support for removing initializer from input to support models converted from old version pytorch where weights are stored as inputs.
- Added a recommended configuration for the IPU_Transformer platform.
- New utilities:
  * Tool for converting the float16 model to the float32 model.
  * Tool for converting the NCHW model to the NHWC model.
  * Tool for quantized models with random input.

  - Three examples for quantization models from Timm, Torchvision, and ONNXRuntime modelzoo respectively.
  - Bugfixes:
    * Fix a bug that weights are quantized with the "NonOverflow" method when using the "MinMSE" method.

- Pytorch Quantizer
  - Support of some operations quantization in quantizer: inplace div, inplace sub
  - Log and document enhancement to emphasize fast-finetune
  - Timm models quantization script example
  - Bug fix for operators: clamp and prelu
  - QAT Support quantization of operations with multiple outputs
  - QAT EOU enhancements: significantly reduces the need for network modifications
  - QAT ONNX exporting enhancements: support more configurations
  - New QAT examples
- TF2 Quantizer
  - Support for Tensorflow 2.11 and 2.12.
  - Support for the 'tf.linalg.matmul' operator.
  - Updated shift_bias constraints for NPU workflow.
  - Support for dumping models containing operations with multiple outputs.
  - Added an example of a sequential model.
  - Bugfixes:
    * Fix a bug that Hardsigmoid and Hardswish are not mapped to DPU without Batch Normalization.
    * Fix a bug when both align_pool and align_concat are used simultaneously.
    * Fix a bug in the sequential model when a layer has multiple consumers.
- TF1 Quantizer
  - Update shift_bias constraints for NPU workflow.
  - Bugfixes:
    * Fix a bug in fast_finetune when the 'input_node' and 'quant_node' are inconsistent.
    * Fix a bug that AddV2 op identified as BiasAdd.
    * Fix a bug when the data type of the concat op is not float.

* Fix a bug in split_large_kernel_pool when the stride is not equal to 1.

ONNXRuntime Execution Provider

- Support new OPs, such as PRelu, ReduceSum, LpNormlization, DepthToSpace(DCR).
- Increase the percentage of model operators performed on the NPU.
- Fixed some issues causing model operators allocation to CPU.
- Improved report summary
- Support the encryption of the VOE cache
- End-2-End Application support on NPU
  - Enable running pre/post/custom ops on NPU, utilizing ONNX feature of E2E extensions.
  - Two examples published for yolov8 and resnet50, in which preprocessing custom op is added and runs on NPU.
- Performance: latency improves by up to 18% and power savings by up to 35% by additionally running preprocessing on NPU apart from inference.
- Multiple NPU overlays support
  - VOE configuration that supports both CNN-centric and GEMM-centric NPU overlays.
  - Increases number of ops that run on NPU, especially for models which have both GEMM and CNN ops.
  - Examples published for use with some of the vision transformer models.

NPU and Compiler

- New operators support
  - Global average pooling with large spatial dimensions
- Single Activation (no fusion with conv2d, e.g., ReLU or single-alpha PReLU)
- Operator support enhancement
  - Enlarge the width dimension support range for depthwise-conv2d
  - Support more generic broadcast for element-wise like operator
  - Support output channel not aligned with 4B GStiling
  - Support Mul and LeakyRelu fusion
  - Concatenation's redundant input elimination
  - Channel Augmentation for conv2d (3x3, stride=2)

- Performance optimization
  - PDI partition refine to reduce the overhead for PDI swap
  - Enabled cost model for some specific models
- Fixed asynchronous error in multiple thread scenario
- Fixed known issue on tanh and transpose-conv2d hang issue

Known Issues

- Support for multiple applications is limited to up to eight
- Windows Studio Effects should be disabled when using the Latency profile. To disable Windows Studio Effects, open **Settings > Bluetooth & devices > Camera**, select your primary camera, and then disable all camera effects.

3.1.14 Version 0.9

Quantizer

- Pytorch Quantizer
  - Dict input/output support for model forward function
  - Keywords argument support for model forward function
  - Matmul subroutine quantization support
  - Support of some operations in quantizer: softmax, div, exp, clamp
  - Support quantization of some non-standard conv2d.
- ONNX Quantizer
  - Add support for Float16 and BFloat16 quantization.
  - Add C++ kernels for customized QuantizeLinear and DequantizeLinaer operations.
  - Support saving quantizer version info to the quantized models' producer field.
  - Support conversion of ReduceMean to AvgPool in NPU workflow.
  - Support conversion of BatchNorm to Conv in NPU workflow.
  - Support optimization of large kernel GlobalAvgPool and AvgPool operations in NPU workflow.
  - Supports hardware constraints check and adjustment of Gemm, Add, and Mul operations in NPU workflow.
  - Supports quantization for LayerNormalization, HardSigmoid, Erf, Div, and Tanh for NPU.

ONNXRuntime Execution Provider

- Support new OPs, such as Conv1d, LayerNorm, Clip, Abs, Unsqueeze, ConvTranspose.
- Support pad and depad based on NPU subgraph's inputs and outputs.
- Support for U8S8 models quantized by ONNX quantizer.
- Improve report summary tools.

NPU and Compiler

- Supported exp/tanh/channel-shuffle/pixel-unshuffle/space2depth
- Performance uplift of xint8 output softmax
- Improve the partition messages for CPU/DPU
- Improve the validation check for some operators
- Accelerate the speed of compiling large models
- Fix the elew/pool/dwc/reshape mismatch issue and fix the stride_slice hang issue
- Fix str_w != str_h issue in Conv

LLM

- Smoothquant for OPT1.3b, 2.7b, 6.7b, 13b models.
- Huggingface Optimum ORT Quantizer for ONNX and Pytorch dynamic quantizer for Pytorch
- Enabled Flash attention v2 for larger prompts as a custom torch.nn.Module
- Enabled all CPU ops in bfloat16 or float32 with Pytorch
- int32 accumulator in AIE (previously int16)
- DynamicQuantLinear op support in ONNX
- Support different compute primitives for prefill/prompt and token phases
- Zero copy of weights shared between different op primitives
- Model saving after quantization and loading at runtime for both Pytorch and ONNX
- Enabled profiling prefill/prompt and token time using local copy of OPT Model with additional timer instrumentation
- Added demo mode script with greedy, stochastic and contrastive search options

ASR

- Support Whipser-tiny
- All GEMMs offloaded to AIE
- Improved compile time
- Improved WER

Known issues

- Flow control OPs including “Loop”, “If”, “Reduce” not supported by VOE
- Resizing OP in ONNX opset 10 or lower is not supported by VOE
- Tensorflow 2.x quantizer supports models within tf.keras.model only
- Running the quantizer Docker in WSL on Ryzen AI laptops might encounter an OOM (Out-of-Memory) issue.
- Running multiple concurrent models using temporal sharing on the 5x4 binary is not supported
- Only batch sizes of 1 are supported
- Only models with the pretrained weights setting = TRUE should be imported
- Launching multiple processes on 4 1x4 binaries can cause hangs, especially when models have many sub-graphs

3.1.15 Version 0.8

Quantizer

- Pytorch Quantizer
  - Pytorch 1.13 and 2.0 support
  - Mixed precision quantization support, supporting float32/float16/bfloat16/intx mixed quantization
  - Support of bit-wise accuracy cross check between quantizer and ONNX-runtime
  - Split and chunk operators were automatically converted to slicing
  - Add support for BFP data type quantization

  - Support of some operations in quantizer: where, less, less_equal, greater, greater_equal, not, and, or, eq, maximum, minimum, sqrt, Elu, Reduction_min, argmin
  - QAT supports training on multiple GPUs
  - QAT supports operations with multiple inputs or outputs
- ONNX Quantizer
  - Provided Python wheel file for installation
  - Support OnnxRuntime 1.15
  - Supports setting input shapes of random data reader
  - Supports random data reader in the dump model function
  - Supports saving the S8S8 model in U8S8 format for NPU
  - Supports simulation of Sigmoid, Swish, Softmax, AvgPool, GlobalAvgPool, ReduceMean and LeakyRelu for NPU
  - Supports node fusions for NPU

ONNXRuntime Execution Provider

- Supports for U8S8 quantized ONNX models
- Improve the function of falling back to CPU EP
- Improve AIE plugin framework
  - Supports LLM Demo
  - Supports Gemm ASR
  - Supports E2E AIE acceleration for Pre/Post ops
  - Improve the easy-of-use for partition and deployment
- Supports models containing subgraphs
- Supports report summary about OP assignment
- Supports report summary about DPU subgraphs falling back to CPU
- Improve log printing and troubleshooting tools.
- Upstreamed to ONNX Runtime Github repo for any data type support and bug fix

NPU and Compiler

- Extended the support range of some operators
  - Larger input size: conv2d, dwc
  - Padding mode: pad

  - Broadcast: add
  - Variant dimension (non-NHWC shape): reshape, transpose, add
- Support new operators, for example, reducemax(min/sum/avg) and argmax(min).
- Enhanced multi-level fusion
- Performance enhancement for some operators
- Add quantization information validation
- Improvement in device partition
  - User friendly message
  - Target-dependency check

Demos

- New Demos link: https://account.amd.com/en/forms/downloads/ryzen-ai-software-platform-xef.html?filename=transformers_2308.zip
  - LLM demo with OPT-1.3B/2.7B/6.7B
  - Automatic speech recognition demo with Whisper-tiny

Known issues

- Flow control OPs including “Loop”, “If”, “Reduce” not supported by VOE
- Resize OP in ONNX opset 10 or lower not supported by VOE
- Tensorflow 2.x quantizer supports models within tf.keras.model only

Running the quantizer Docker in WSL on Ryzen AI laptops might encounter an OOM (Out-of-Memory) issue. - Run multiple concurrent models by temporal sharing on the Performance optimized overlay (5x4.xclbin) is not supported - Support batch size 1 only for NPU

3.1.16 Version 0.7

Quantizer

- Docker Containers
  - Provided CPU dockers for Pytorch, Tensorflow 1.x, and Tensorflow 2.x quantizer
  - Provided GPU Docker files to build GPU dockers

- Pytorch Quantizer
  - Supports multiple output conversion to slicing
  - Enhanced transpose OP optimization
  - Inspector support new IP targets for NPU
- ONNX Quantizer
  - Provided Python wheel file for installation
  - Supports quantizing ONNX models for NPU as a plugin for the ONNX Runtime native quantizer
  - Supports power-of-two quantization with both QDQ and QOP format
  - Supports Non-overflow and Min-MSE quantization methods
  - Supports various quantization configurations in power-of-two quantization in both QDQ and QOP format.
    - Supports signed and unsigned configurations.
    - Supports symmetry and asymmetry configurations.
    - Supports per-tensor and per-channel configurations.
  - Supports bias quantization using int8 datatype for NPU.
  - Supports quantization parameters (scale) refinement for NPU.
  - Supports excluding certain operations from quantization for NPU.
  - Supports ONNX models larger than 2 GB.
  - Supports using CUDAExecutionProvider for calibration in quantization
  - Open source and upstreamed to Microsoft Olive Github repo
- TensorFlow 2.x Quantizer
  - Added support for exporting the quantized model ONNX format.
  - Added support for the keras.layers.Activation('leaky_relu')
- TensorFlow 1.x Quantizer
  - Added support for folding Reshape and ResizeNearestNeighbor operators.
  - Added support for splitting Avgpool and Maxpool with large kernel sizes into smaller kernel sizes.
  - Added support for quantizing Sum, StridedSlice, and Maximum operators.
  - Added support for setting the input shape of the model, which is useful in deploying models with undefined input shapes.

  - Add support for setting the opset version in exporting ONNX format

ONNX Runtime Execution Provider

- Vitis ONNX Runtime Execution Provider (VOE)
  - Supports ONNX Opset version 18, ONNX Runtime 1.16.0, and ONNX version 1.13
  - Supports both C++ and Python APIs(Python version 3)
  - Supports deploy model with other EPs
  - Supports falling back to CPU EP
  - Open source and upstreamed to ONNX Runtime Github repo
  - Compiler
    * Multiple Level op fusion
    * Supports the same muti-output operator like chunk split
    * Supports split big pooling to small pooling
    * Supports 2-channel writeback feature for Hard-Sigmoid and Depthwise-Convolution
    * Supports 1-channel GStiling
    * Explicit pad-fix in CPU subgraph for 4-byte alignment
    * Tuning the performance for multiple models

NPU

- Two configurations
  - Power Optimized Overlay
    * Suitable for smaller AI models (1x4.xclbin)
    * Supports spatial sharing, up to 4 concurrent AI workloads
  - Performance Optimized Overlay (5x4.xclbin)
    * Suitable for larger AI models

Known issues

- Flow control OPs including “Loop”, “If”, “Reduce” are not supported by VOE
- Resize OP in ONNX opset 10 or lower not supported by VOE
- Tensorflow 2.x quantizer supports models within tf.keras.model only

- Running the quantizer Docker in WSL on Ryzen AI laptops might encounter an OOM (Out-of-Memory) issue.
- Run multiple concurrent models by temporal sharing on the Performance optimized overlay (5x4.xclbin) is not supported

3.2 Installation Instructions

This page covers Ryzen AI installation on Windows.

3.2.1 Prerequisites

The Ryzen AI Software supports AMD processors with a Neural Processing Unit (NPU). Refer to the release notes for the full list of *supported configurations*.

The following dependencies must be installed on the system before installing the Ryzen AI Software:

Dependencies	Version Requirement
Windows 11	build >= 22621.3527
Visual Studio	2022
cmake	version >= 3.26
Python distribution (Miniforge preferred)	Latest version

IMPORTANT:

- Visual Studio 2022 Community (Optional for AMD Quark, to support custom op flow): ensure that *Desktop Development with C++* is installed
- Miniforge: ensure that the following path is set in the System PATH variable: `path\to\miniforge3\condabin` or `path\to\miniforge3\Scripts\` or `path\to\miniforge3\` (The System PATH variable should be set in the *System Variables* section of the *Environment Variables* window).

3.2.2 Install NPU Drivers

- Download and Install the NPU driver version: 32.0.203.280 or newer using the following links:
  - `NPU Driver (Version 32.0.203.280)`
    - NPU driver 32.0.203.376 is production driver for Phoenix, Hawk Point, Strix, Strix Halo, and Krackan Point.
  - `NPU Driver (Version 32.0.203.376)`
- Install the NPU drivers by following these steps:
  - Extract the downloaded ZIP file.
  - Open a terminal in administrator mode and execute the `.\npu_sw_installer.exe` file.
- Ensure that NPU driver (Version:32.0.203.280, Date:5/16/2025) is correctly installed by opening Task Manager -> Performance -> NPU0.

3.2.3 Install Ryzen AI Software

- Download the Ryzen AI Software installer `ryzen-ai-1.8.0.exe`.
- Launch the EXE installer and follow the instructions on the installation wizard:
  - Accept the terms of the Licence agreement
  - Provide the destination folder for Ryzen AI installation (default: `C:\Program Files\RyzenAI\1.8.0`)
  - Specify the name for the conda environment (default: `ryzen-ai-1.8.0`)

The Ryzen AI Software packages are now installed in the conda environment created by the installer.

> **Note**
>
> NuGet package is available to download at `ryzen_ai_nuget_1.8.0.zip`.

3.2.4 Test the Installation

The Ryzen AI Software installation folder contains test to verify that the software is correctly installed. This installation test can be found in the `quicktest` subfolder which is expected to work for Strix (STX) or newer devices.

- Open a Conda command prompt (search for "Miniforge Prompt" in the Windows start menu)
- Activate the Conda environment created by the Ryzen AI installer:

```
conda activate ryzen-ai-<version>
```

- Run the test:

```
cd %RYZEN_AI_INSTALLATION_PATH%/quicktest
python quicktest.py
```

```
[I:onnxruntime:, session_state_utils.cc:243 onnxruntime::session_state_
↪utils::SaveInitializedTensors] Saving initialized tensors.
[I:onnxruntime:, session_state_utils.cc:438 onnxruntime::session_state_
↪utils::SaveInitializedTensors] Done saving initialized tensors
[I:onnxruntime:, inference_session.cc:2532␣
↪onnxruntime::InferenceSession::Initialize] Session successfully␣
↪initialized.
Test Finished
```

- Verify NPU activity by opening **Task Manager** → **Performance** → **NPU** while the test is running. You should see NPU utilization increase during model inference.

To enable NPU offloading logs using *ONNX Runtime session options*, modify the `quicktest.py` script to include session options that control logging verbosity. The following code snippet demonstrates how to set up the session options for detailed logging:

```
# Create session options
session_options = ort.SessionOptions()
session_options.log_severity_level = 0  # 0=Verbose, 1=Info, 2=Warning,␣
↪3=Error, 4=Fatal

try:
    session = ort.InferenceSession(model,
                                   sess_options=session_options,
                                   providers=providers,
                                   provider_options=provider_options)
except Exception as e:
    print(f"Failed to create an InferenceSession: {e}")
    sys.exit(1)  # Exit the program with a non-zero status to indicate an␣
↪error
```

- Run the updated `quicktest.py` script with logs filter:

```
cd %RYZEN_AI_INSTALLATION_PATH%/quicktest
python quicktest.py 2>&1 | findstr /i "Vitis | NPU | Test"
```

- On a successful run, the output will look ssimilar to the one shown below. This indicates that the model is running on the NPU and that the installation of the Ryzen AI Software was

successful:

```
[I:onnxruntime:, stat.cpp:198 stat.cpp] [Vitis AI EP] No. of Operators :
[I:onnxruntime:, stat.cpp:198 stat.cpp]    NPU   398
[I:onnxruntime:, stat.cpp:198 stat.cpp] VITIS_EP_CPU     2
[I:onnxruntime:, stat.cpp:198 stat.cpp] [Vitis AI EP] No. of Subgraphs :
[I:onnxruntime:, stat.cpp:198 stat.cpp]    NPU     1
[I:onnxruntime:, stat.cpp:198 stat.cpp] Actually running on NPU      1
[I:onnxruntime:, vitisai_compile_model.cpp:1488 vitisai_compile_model.
↪cpp] AVG CPU Usage 4.93969%
[I:onnxruntime:, vitisai_compile_model.cpp:1489 vitisai_compile_model.
↪cpp] Peak Working Set size 565.277 MB
[V:onnxruntime:, session_state.cc:1350␣
↪onnxruntime::VerifyEachNodeIsAssignedToAnEp]  All nodes placed on␣
↪[VitisAIExecutionProvider]. Number of nodes: 3
[I:onnxruntime:, unzip.hpp:18 unzip.hpp] inflateInit successful,␣
↪available input: 137827
Test finished
```

Note

- The full path to the Ryzen AI Software installation folder is stored in the `RYZEN_AI_INSTALLATION_PATH` environment variable.

3.3 Examples, Demos, Tutorials

This page introduces various demos, examples, and tutorials currently available with the Ryzen™ AI Software.

3.3.1 Getting Started Tutorials

NPU

- Getting Started Tutorial for INT8 models - Uses a custom ResNet model to demonstrate:
  - Pretrained model conversion to ONNX
  - Model Quantization using AMD Quark quantizer
  - Deployment using ONNX Runtime C++ and Python code
- Getting Started Tutorial for BF16 models - Uses a custom ResNet model to demonstrate:
  - Preparation and compilation of BF16 models
  - Deployment using Python

    - Deployment using C++
- Hello World Jupyter Notebook Tutorial
- Additional BF16 model examples:
    - Image Classification
    - Finetuned DistilBERT for Text Classification
- Super-Resolution Models on AMD Ryzen AI NPU
- LLM examples
    - LLMs on RyzenAI with ONNX Runtime GenAI API
    - ONNX Runtime GenAI(OGA)-based RAG LLM
    - Running Vision Language Model (VLM) on RyzenAI NPU
    - Running GPT-OSS-20B with chat template

iGPU

- ResNet50 on iGPU

3.3.2 Other examples, demos, and tutorials

- Refer to RyzenAI-SW repo

3.4 Linux Installation Instructions

Ryzen AI for Linux supports running AI models on the AMD Neural Processing Unit (NPU). Linux support is available on newer consumer APUs, beginning with Strix Point, and on newer embedded APUs, beginning with Ryzen™ AI Embedded P100.

With this release, Users can now compile and run AI models using the following formats:

- CNN models in INT8
- CNN models in BF16
- NLP models (e.g., BERT, encoder-based) in BF16
- LLMs (NPU-only flow)

3.4.1 Prerequisites

Use the commands below to install Python 3.12.x along with certain dependencies

```
sudo apt update
sudo apt install python3.12
sudo apt install python3.12-venv
sudo apt install libboost-filesystem1.74.0
sudo apt install dkms
```

After installing required Ubuntu distribution and Python version, proceed with NPU drivers installation

3.4.2 Install NPU Drivers

- Download the NPU driver package from *Downloads* section of Ryzen AI Software Drivers.
- **RyzenAI linux driver package contains**
  - **XRT Package**
    - xrt_202620.2.25.37_24.04-amd64-base.deb
    - xrt_202620.2.25.37_24.04-amd64-base-dev.deb
    - xrt_202620.2.25.37_24.04-amd64-npu.deb
  - **NPU driver package**
    - xrt_plugin.2.25.260102.56.release_24.04-amd64-amdxdna.deb
- Install NPU driver package on your machine

```
sudo apt install --fix-broken -y ./xrt_202620.2.25.37_24.04-amd64-base.deb
sudo apt install --fix-broken -y ./xrt_202620.2.25.37_24.04-amd64-base-
↪dev.deb
sudo apt install --fix-broken -y ./xrt_202620.2.25.37_24.04-amd64-npu.deb
sudo apt install --fix-broken -y ./xrt_plugin.2.25.260102.56.release_24.
↪04-amd64-amdxdna.deb
```

- Verify your Driver installation

```
 source /opt/xilinx/xrt/setup.sh
 xrt-smi examine

Device(s) Present
|BDF             |Name                |Architecture  |Topology  |
|----------------|--------------------|--------------|----------|
|[0000:c5:00.1]  |NPU Strix           |aie2p         |6x8       |

 # NPU name and Device code may differ based on your machine
```

3.4.3 Install Ryzen AI Software

- Download the RyzenAI for Linux package *ryzen_ai-1.8.0.tgz* from *Downloads* section of Ryzen AI Software Installer.
- Navigate to the downloaded path and follow the below steps

```
mkdir ryzen_ai-1.8.0
cp ryzen_ai-1.8.0.tgz ryzen_ai-1.8.0

cd ryzen_ai-1.8.0
tar -xvzf ryzen_ai-1.8.0.tgz
```

- Install RyzenAI package at your desired target path

```
./install_ryzen_ai.sh -a yes -p <TARGET-PATH>/venv
source <TARGET-PATH>/venv/bin/activate
```

- This will successfully install RyzenAI and activate the Virtual environment at your target location

```
# Validate your installation path
echo $RYZEN_AI_INSTALLATION_PATH
```

- Set essential Environment variables

```
export LD_LIBRARY_PATH=/lib/x86_64-linux-gnu:${RYZEN_AI_INSTALLATION_PATH}
↪/onnxruntime/lib/:$LD_LIBRARY_PATH
```

- Important Note

```
# After activating the Ryzen AI Virtual environment-
  source <TARGET-PATH>/venv/bin/activate

# Always run the below command. It appends necessary XRT utils to LD_
↪LIBRARY_PATH for successful execution on NPU
  source /opt/xilinx/xrt/setup.sh
```

3.4.4 Test the Installation

The RyzenAI software package contains a test script that verifies your correct installation of NPU Drivers.

- Navigate to your targeted Virtual Environment created in the previous step
- You will observe a subfolder named “quicktest”

```
cd <TARGET-PATH>/venv/quicktest
python quicktest.py
```

- The quicktest.py script picks up a simple CNN model, compiles it and runs on AMD's Neural Processing Unit (NPU).
- On successful run, you can observe output as shown below.

```
Setting environment for STX/KRK

Test Finished
```

3.4.5 Examples, Demos, Tutorials

- RyzenAI-SW demonstrates various demos and examples for Model compilation and deployment on NPUs
- **Here are a few examples from our RyzenAI Software Repository**
  - Getting started Resnet with BF16 Model
  - Getting started Resnet with INT8 Model
  - Yolov8m Model for Object Detection

3.4.6 Note

Before running the above examples -

- RyzenAI-SW repo hosts a diverse set of examples with both NPU and iGPU Execution Provider. However, Linux currently supports NPU only flow.
- RyzenAI creates its own Python Virtual Environment to run the examples. Skip conda environment instruction as they are Windows specific only
- Ensure to activate Linux based Python Virtual Environment and load necessary XRT Utils

```
source <TARGET-PATH>/venv/bin/activate

source /opt/xilinx/xrt/setup.sh
```

Get NPU Info for your Machine

Getting started Resnet with INT8 Model

Getting started Resnet with INT8 Model contains Resnet_util.py script that has a function "get_npu_info" to detect correct "NPU type" in your machine. This NPU lookup logic is based for Windows system.

For Linux, NPU lookup logic is shown below:

```
import subprocess

def get_npu_info():
    # Run below command as subprocess to enumerate PCI devices
    command = r'lspci -nn'
    process = subprocess.Popen(command, shell=True, stdout=subprocess.
↪PIPE, stderr=subprocess.PIPE)
    stdout, stderr = process.communicate()

    # Check for supported Hardware IDs
    npu_type = ''
    if '1022:17f0' in stdout.decode(): npu_type = 'STX/KRK'
    return npu_type
```

3.4.7 Running LLM

Follow this page to run LLM models on Linux: llm_linux

3.5 Model Quantization

Model quantization is the process of mapping high-precision weights/activations to a lower precision format, such as BF16/INT8, while maintaining model accuracy. This technique enhances the computational and memory efficiency of the model for deployment on NPU devices. It can be applied post-training, allowing existing models to be optimized without the need for retraining.

The Ryzen AI compiler supports input models in the following formats:

- CNN Models
  - INT8 (quantized)
  - FP32 (automatically converted to BF16 during compilation)
- Transformer Models:
  - FP32 (automatically converted to BF16 during compilation)

Ryzen AI Software natively supports both CNN and Transformer models in floating-point (FP32) format. When FP32 models are provided as input, the VitisAI EP automatically converts them to bfloat16 (BF16) precision and processes them through the optimized BF16 compilation pipeline.

For CNN models, AMD Quark quantization enables conversion to INT8 format, delivering improved inference performance compared to higher precision alternatives. This quantization pathway provides an additional optimization option for CNN workloads requiring maximum efficiency.

The complete list of operations supported for different quantization types can be found in *Supported Operations*.

3.5.1 FP32 to BF16 Conversion

Ryzen AI provides seamless support for deploying original floating-point (FP32) models on NPU hardware through automatic conversion to BFLOAT16 (BF16) format. The conversion from FP32 to BF16 is performed when the model is compiled by the VitisAI EP. BF16 is a 16-bit floating-point format designed to have the same exponent size as FP32, allowing a wide dynamic range, but with reduced precision to save memory and speed up computations. This feature enables developers to deploy models in their native format while leveraging the Ryzen AI automatic conversion for efficient execution on NPU.

FP32 to BF16 Examples

Explore these practical examples demonstrating FP32 to BF16 conversion across different CNN, NLP model types:

- Image Classification using ResNet50 model on NPU
- Finetuned DistilBERT for Text Classification
- Advanced quantization techniques Fast Finetuning for BF16 models.

3.5.2 FP32 to INT8 Conversion

Quantization to INT8 format introduces several challenges, primarily revolving around the potential drop in model accuracy. Choosing the right quantization parameters—such as data type, bit-width, scaling factors, and the decision between per-channel or per-tensor quantization—adds layers of complexity to the design process. These decisions significantly impact both model accuracy and performance. While **AMD Quark** is the recommended quantization tool, third-party tools that support QDQ (Quantize-Dequantize) operations can also be used for model quantization.

RyzenAI supports the following INT8 datatypes:

- XINT8: uses symmetric INT8 activation and weights quantization with power-of-two scales
- A8W8: uses symmetric INT8 activation and weights quantization with float scales
- A16W8: uses symmetric INT16 activation and symmetric INT8 weights quantization with float scales

AMD Quark is the recommended quantization tool to convert FP32 models to INT8. But third-party tools that support QDQ (Quantize-Dequantize) operations can also be used for model quantization to A8W8 and A16W8.

AMD Quark

AMD Quark is a comprehensive cross-platform deep learning toolkit designed to simplify and enhance the quantization of deep learning models. Supporting both PyTorch and ONNX models, Quark empowers developers to optimize their models for deployment on a wide range of hardware backends, achieving significant performance gains without compromising accuracy.

AMD Quark provides default configurations that support INT8 quantization. For example, *XINT8* uses symmetric INT8 activation and weights quantization with power-of-two scales using the MinMSE calibration method.

For more challenging model quantization needs, **AMD Quark** supports different quantization configurations such as *A8W8*, *A16W8*, and advanced quantization techniques. For more details, refer to AMD Quark for Ryzen AI

The quantization configuration can be customized using the *QuantizationConfig* class. The following example shows how to set up the quantization configuration for INT8 quantization:

```
quant_config = QuantizationConfig(calibrate_method=PowerOfTwoMethod.
↪MinMSE,
                                  activation_type=QuantType.QUInt8,
                                  weight_type=QuantType.QInt8,
                                  enable_npu_cnn=True,
                                  extra_options={'ActivationSymmetric':␣
↪True})
config = Config(global_quant_config=quant_config)
print("The configuration of the quantization is {}".format(config))
```

The user can use the *get_default_config('XINT8')* function to get the default configuration for INT8 quantization.

FP32 to INT8 Examples

Explore practical INT8 quantization examples:

- Running INT8 model on NPU using Getting Started Tutorial
- AMD Quark Tutorial for Ryzen AI Deployment
- Advanced quantization techniques Fast Finetuning and Cross Layer Equalization for INT8 model

3.6 Model Compilation and Deployment

3.6.1 Introduction

The Ryzen AI Software supports models saved in the ONNX format and uses ONNX Runtime as the primary mechanism to load, compile and run models.

NOTE: Models with ONNX opset 17 are recommended. If your model uses a different opset version, consider converting it using the ONNX Version Converter

For a complete list of supported operators, consult this page: *Supported Operators*.

Loading Models

Models are loaded by creating an ONNX Runtime `InferenceSession` using the Vitis AI Execution Provider (VAI EP):

```
import onnxruntime

session_options = onnxruntime.SessionOptions()
vai_ep_options  = {}                            # Vitis AI EP options go␣
↪here

session = onnxruntime.InferenceSession(
    path_or_bytes = model,                      # Path to the ONNX model
    sess_options = session_options,             # Standard ORT options
    providers = ['VitisAIExecutionProvider'], # Use the Vitis AI␣
↪Execution Provider
    provider_options = [vai_ep_options]         # Pass options to the Vitis␣
↪AI Execution Provider
)
```

The `provider_options` parameter enables the configuration of the Vitis AI Execution Provider (EP). For a comprehensive list of supported provider options, refer to the *Vitis AI EP Options Reference Guide* below.

When a model is first loaded into an ONNX Runtime (ORT) inference session, it is compiled into the format required by the NPU. The resulting compiled output can be saved as an ORT EP context file or stored in the Vitis AI EP cache directory.

If a compiled version of the ONNX model is already available — either as an EP context file or within the Vitis AI EP cache — the model will not be recompiled. Instead, the precompiled version will be loaded automatically. This greatly reduces session creation time and improves overall efficiency. For more details, refer to the section on *Managing Compiled Models*.

Deploying Models

Once the ONNX Runtime inference session is initialized and the model is compiled, the model is deployed using the ONNX Runtime `run()` API:

```
input_data = {}
for input in session.get_inputs():
    input_data[input.name] = ...  # Initialize input tensors
```

(continues on next page)

(continued from previous page)

```
outputs = session.run(None, input_data) # Run the model
```

The ONNX graph is automatically partitioned into multiple subgraphs by the Vitis AI Execution Provider (EP). During deployment, the subgraph(s) containing operators supported by the NPU are executed on the NPU. The remaining subgraph(s) are executed on the CPU. This graph partitioning and deployment technique across CPU and NPU is fully automated by the VAI EP and is totally transparent to the end-user.

3.6.2 Vitis AI EP Options Reference Guide

VitisAI EP Provider Options

The `provider_options` parameter of the ORT `InferenceSession` allows passing options to configure the Vitis AI EP. The following options are supported:

Table 2: Vitis AI EP Provider Options

Option	Description	Values / Type	Default
config_fi	Config file for BF16 compilation options. See *Config File Options*.	String	N/A
target	Set which Vitis AI EP backend to use for compiling/running integer model. For details see *Using INT8 Models*.	X2, X1	X2
xclbin	To be used only when running INT8 CNN models on PHX/HPT devices. For details see *Using INT8 Models*	String	None
encryptio	256-bit key for encrypting EP context model. See *EP Context Cache*.	String (64 hex)	None
opt_level	Compiler optimization level for INT8 only.	0, 1, 2, 3, 65536 (maximum effort, experimental)	0
cache_dir	VitisAI cache directory. Set enable_cache_file_io_in_mem to 0.	String	C:\temp\%USERNAME%\vaip\.cache
cache_key	Subfolder in cache for compiled model. Set enable_cache_file_io_in_mem to 0.	String	MD5 hash of model
enable_ca	Keep compiled model in memory (1) or save to disk (0). • 0: saves to disk (BF16 and INT8 models) • 1: saves in memory (INT8 models only)	0, 1	1
ai_analyz	Enable compile-time analysis data.	Boolean	False
ai_analyz	Enable inference-time analysis data.	Boolean	False

Enabling ONNX Runtime Logs

To enable detailed logging for debugging purposes, set the ONNX Runtime session log severity level using `SessionOptions.log_severity_level`:

```
import onnxruntime

session_options = onnxruntime.SessionOptions()
session_options.log_severity_level = 1  # Set log level (see table below)

vai_ep_options = {}
```

(continues on next page)

(continued from previous page)

```
session = onnxruntime.InferenceSession(
    path_or_bytes = model,
    sess_options = session_options,
    providers = ['VitisAIExecutionProvider'],
    provider_options = [vai_ep_options]
)
```

Table 3: ORT Log Severity Levels

Level	Description	Value
Verbose	All messages (most detailed)	0
Info	Informational messages and above	1
Warning (Default)	Warnings and above	2
Error	Errors and above	3
Fatal	Fatal errors only	4

NOTE: The `log_level` parameter in provider options has been deprecated. To enable logging, use `SessionOptions.log_severity_level` as shown in the example above.

Config File Options

When compiling BF16 models, a JSON configuration file can be provided to the VitisAI EP using the `config_file` provider option. This configuration file is used to specify additional options to the compiler.

The default the configuration file for compiling BF16 models contains the following:

```
{
 "passes": [
     {
         "name": "init",
         "plugin": "vaip-pass_init"
     },
     {
         "name": "vaiml_partition",
         "plugin": "vaip-pass_vaiml_partition",
         "vaiml_config": {
             "optimize_level": 1,
             "preferred_data_storage": "auto"
         }
     }
 ],
```

(continues on next page)

(continued from previous page)

```
 "target": "VAIML",
    "targets": [
        {
            "name": "VAIML",
            "pass": [
                "init",
                "vaiml_partition"
            ]
        }
    ]
}
```

The `vaiml_config` section of the configuration file contains the user options. The supported user options are described below.

Table 4: Config File Options (vaiml_config)

Option	Description	Values	Default
`optimize_leve`	Compiler optimization level.	1, 2, 3	1
`preferred_dat`	Data layout: "auto" (let compiler choose), "vectorized" (for CNNs), "unvectorized" (for Transformers).	auto, vectorized, unvectorized	auto

3.6.3 Using BF16 Models

When compiling BF16 models, an optional configuration file can be provided to the VitisAI EP. This file is specified using the `config_file` provider option. For more details, refer to *Config File Options* section.

NOTE:

- Running BF16 Models is only supported for STX/KRK or newer devices. For the model compatibility table see *Release Notes*.
- For C++ applications that need to compile BF16 models at runtime, include `${CONDA_PREFIX}/Lib/site-packages/flexml/flexml_extras/lib/vaiml.dll` along with the DLLs specified in Application Development. However, using pre-compiled BF16 models is recommended for C++ deployment.

Sample Python Code

Python example loading a configuration file called vai_ep_config.json:

```
import onnxruntime

vai_ep_options = {
    'config_file': 'vai_ep_config.json',
}

session = onnxruntime.InferenceSession(
    "resnet50.onnx",
    providers=['VitisAIExecutionProvider'],
    provider_options=[vai_ep_options]
)
```

Sample C++ Code

C++ example loading a configuration file called vai_ep_config.json:

```
#include <onnxruntime_cxx_api.h>

auto onnx_model = "resnet50.onnx"
Ort::Env env(ORT_LOGGING_LEVEL_WARNING, "resnet50_bf16");
auto session_options = Ort::SessionOptions();
auto vai_ep_options = std::unorderd_map<std::string,std::string>({});
vai_ep_options["config_file"] = "vai_ep_config.json";
session_options.AppendExecutionProvider_VitisAI(vai_ep_options);
auto session = Ort::Session(
    env,
    std::basic_string<ORTCHAR_T>(onnx_model.begin(), onnx_model.end()).c_
↪str(),
    session_options);
```

3.6.4 Using INT8 Models

Ryzen AI features a new compiler for INT8 models. This compiler is enabled by default and provides the following improvements:

- Improved ease of use and enhanced performance for models running on STX, KRK, and later devices.
- General asymmetric quantization support to enable third-party quantized models
- Support for XINT8, A8W8, and A16W8 quantization configuration providing greater flexibility for model optimization.

The `target` provider options can be used to select which backend to use when compiling the INT8 model. The option accepts the following values:

- `X2` — Default backend for integer models. Supports STX, KRK and newer devices.
- `X1` — Legacy backend for integer models. Supports PHX, HPT, STX and KRK devices. This setting should be used when running on PHX and HPT devices. It can also be used on STX and KRK devices in the cases where better results are achieved than with the default X2 setting.

Device-Specific Settings

Suitable settings for the `target` and `xclbin` provider options are dependent on the type of device. The application must perform a device detection check before configuring the Vitis AI EP. For more details on how to do this, refer to the *Application Development* page.

When compiling INT8 models on STX/KRK devices:

- The `target` provider option can be set to `X2` (default) or `X1` (a legacy backend might provide better results for some models).
- The `xclbin` provider option must not be set.

When compiling INT8 models on PHX/HPT devices:

- The `target` provider option must be set to `X1`.
- The `xclbin` provider option must be set to `%RYZEN_AI_INSTALLATION_PATH%\voe-4.0-win_amd64\xclbins\phoenix\4x4.xclbin` or to a copy of this file included in the final version of the application. The legacy "1x4" and "Nx4" xclbin files are no longer supported and should not be used.

Sample Python Code

Python example code for running an INT8 model on STX/KRK NPU (target=X2, no xclbin):

```
import os
import onnxruntime

vai_ep_options = {
    'cache_dir': str(cache_dir),
    'cache_key': 'resnet_trained_for_cifar10',
    'enable_cache_file_io_in_mem':'0',
    'target': 'X2' # Default option 'X2'
}

session = onnxruntime.InferenceSession(
    "resnet50_int8.onnx",
    providers=['VitisAIExecutionProvider'],
    provider_options=[vai_ep_options]
)
```

Sample C++ Code

C++ example code for running an INT8 model on STX/KRK NPU (target=X2, no xclbin):

```
#include <onnxruntime_cxx_api.h>

auto onnx_model = "resnet50_int8.onnx"
Ort::Env env(ORT_LOGGING_LEVEL_WARNING, "resnet50_int8");
auto session_options = Ort::SessionOptions();
auto vai_ep_options = std::unordered_map<std::string,std::string>({});
vai_ep_options["cache_dir"]   = exe_dir + "\\my_cache_dir";
vai_ep_options["cache_key"]   = "resnet_trained_for_cifar10";
vai_ep_options["enable_cache_file_io_in_mem"]   = "0";
vai_ep_options["target"]   = "X2";
session_options.AppendExecutionProvider_VitisAI(vai_ep_options);
auto session = Ort::Session(
    env,
    std::basic_string<ORTCHAR_T>(onnx_model.begin(), onnx_model.end()).c_
↪str(),
    session_options);
```

3.6.5 Managing Compiled Models

To avoid the overhead of recompiling models, it is very advantageous to save the compiled models and use these pre-compiled versions in the final application. Pre-compiled models can be loaded instantaneously and immediately executed on the NPU. This greatly improves the session creation time and overall end-user experience.

The RyzenAI Software supports two mechanisms for saving and reloading compiled models:

- VitisAI EP Cache
- ONNX Runtime EP Context Cache

TIP: The VitisAI EP Cache mechanism is most convenient to quickly iterate during the development cycle. The OnnxRuntime EP Context Cache mechanism is recommended for the final version of the application.

VitisAI EP Cache

The VitisAI EP includes a built-in caching mechanism. When a model is compiled for the first time, it is automatically saved in the VitisAI EP cache directory. Any subsequent creation of an ONNX Runtime session using the same model will load the precompiled model from the cache directory, thereby reducing session creation time.

The VitisAI EP Cache mechanism can be used to quickly iterate during the development cycle, but it is not recommended for the final version of the application.

Cache directories generated by the Vitis AI Execution Provider should not be reused across different versions of the Vitis AI EP or across different version of the NPU drivers.

If using the VitisAI EP Cache the application should check the version of the Vitis AI EP and of the NPU drivers. If the application detects a version change, it should delete the cache, or create a new cache directory with a different name.

The location of the VitisAI EP cache is specified with the `cache_dir` and `cache_key` provider options. For INT8 models, the `enable_cache_file_io_in_mem` must be set to 0 otherwise the output of the compiler is kept in memory and is not saved to disk.

Python example:

```
import onnxruntime
from pathlib import Path

vai_ep_options = {
    'cache_dir': str(Path(__file__).parent.resolve()),
    'cache_key': 'compiled_resnet50_int8',
    'enable_cache_file_io_in_mem': 0
}

session = onnxruntime.InferenceSession(
    "resnet50_int8.onnx",
    providers=['VitisAIExecutionProvider'],
    provider_options=[vai_ep_options]
)
```

In the example above, the cache directory is set to the absolute path of the folder containing the script being executed. Once the session is created, the compiled model is saved inside a subdirectory named `compiled_resnet50_int8` within the specified cache folder.

ONNX Runtime EP Context Cache

The Vitis AI EP supports the ONNX Runtime EP context cache feature. This features allows dumping and reloading a snapshot of the EP context before deployment.

The user can enable dumping of the EP context by setting the `ep.context_enable` session option to 1.

The following options can be used for additional control:

- `ep.context_file_path` – Specifies the output path for the dumped context model.
- `ep.context_embed_mode` – Embeds the EP context into the ONNX model when set to 1.

For further details, refer to the official ONNX Runtime documentation: https://onnxruntime.ai/docs/execution-providers/EP-Context-Design.html

EP Context Encryption

By default, the generated context model is unencrypted and can be used directly during inference. If needed, the context model can be encrypted using one of the methods described below.

User-managed encryption

After the context model is generated, the developer can encrypt the generated file using a method of choice. At runtime, the encrypted file can be loaded by the application, decrypted in memory and passed as a serialized string to the inference session.

This method gives complete control to the developer over the encryption process.

EP-managed encryption

The VitisAI EP can optionally encrypt the EP context model using AES256. This is enabled by passing an encryption key using the `encryption_key` VAI EP provider options. The key is a 256-bit value represented as a 64-digit string. At runtime, the same encryption key must be provided to decrypt and load the context model.

With this method, encryption and decryption is seamlessly managed by the VitisAI EP.

Python example:

```
import onnxruntime

vai_ep_options = {
    'encryptionKey':
↪'89703f950ed9f738d956f6769d7e45a385d3c988ca753838b5afbc569ebf35b2'
}

# Compilation session
session_options = ort.SessionOptions()
session_options.add_session_config_entry('ep.context_enable', '1')
session_options.add_session_config_entry('ep.context_file_path', 'context_
↪model.onnx')
session_options.add_session_config_entry('ep.context_embed_mode', '1')
session = ort.InferenceSession(
    path_or_bytes='resnet50_int8.onnx',  # Load the ONNX model
    sess_options=session_options,
    providers=['VitisAIExecutionProvider'],
    provider_options=[vai_ep_options]
)

# Inference session
session_options = ort.SessionOptions()
```

(continues on next page)

(continued from previous page)

```
session = ort.InferenceSession(
    path_or_bytes='context_model.onnx', # Load the EP context model
    sess_options=session_options,
    providers=['VitisAIExecutionProvider'],
    provider_options=[vai_ep_options]
)
```

C++ example:

```
Ort::Env env(ORT_LOGGING_LEVEL_WARNING, "ort");

// VAI EP Provider options
auto vai_ep_options = std::unorderd_map<std::string,std::string>({});
vai_ep_options["encryption_key"] =
 ↪"89703f950ed9f738d956f6769d7e45a385d3c988ca753838b5afbc569ebf35b2";

// Session options
auto session_options = Ort::SessionOptions();
session_options.AppendExecutionProvider_VitisAI(vai_ep_options);

// Inference session
auto onnx_model = "context_model.onnx"; // The EP context model
auto session = Ort::Session(
    env,
    std::basic_string<ORTCHAR_T>(onnx_model.begin(), onnx_model.end()).c_
 ↪str(),
    session_options);
```

NOTE: It is possible to precompile the EP context model using Python and to deploy it using a C++ program.

3.6.6 Operator Assignment Report

The compiler can optionally generate a report on operator assignments across CPU and NPU. To generate this report:

- The `enable_cache_file_io_in_mem` provider option must be set to 0
- The XLNX_ONNX_EP_REPORT_FILE environment variable must be used to specify the name of the generated report. For instance:

```
set XLNX_ONNX_EP_REPORT_FILE=vitisai_ep_report.json
```

When these conditions are satisfied, the report file is automatically generated in the cache directory. This report includes information such as the total number of nodes, the list of operator types in the model, and which nodes and operators run on the NPU or on the CPU. Additionally, the report includes node statistics, such as input to a node, the applied operation, and output from the node.

When these conditions are satisified, the report file is automatically generated in the cache directory. This report includes information such as the total number of nodes, the list of operator types in the model, and which nodes and operators runs on the NPU or on the CPU. Additionally, the report includes node statistics, such as input to a node, the applied operation, and output from the node.

```
{
 "deviceStat": [
 {
  "name": "all",
  "nodeNum": 400,
  "supportedOpType": [
   "::Add",
   "::Conv",
   ...
  ]
 },
 {
  "name": "CPU",
  "nodeNum": 2,
  "supportedOpType": [
   "::DequantizeLinear",
   "::QuantizeLinear"
  ]
 },
 {
  "name": "NPU",
  "nodeNum": 398,
  "supportedOpType": [
   "::Add",
   "::Conv",
   ...
  ]
  ...
```

To disable generation of the report, unset the XLNX_ONNX_EP_REPORT_FILE environment variable:

```
set XLNX_ONNX_EP_REPORT_FILE=
```

3.7 Application Development

This page captures requirements and recommendations for developers looking to create, package and distribute applications targeting NPU-enabled AMD processors.

3.7.1 VitisAI EP / NPU Driver Compatibility

For each version of the VitisAI EP, the EP bounds compatible drivers by a minimum version and a maximum release date.

The following table summarizes the driver requirements for the different versions of the VitisAI EP.

VitisAI EP version	Minimum NPU Driver version	Maximum NPU Driver release date
1.7	32.0.203.280	Jan 22nd, 2029
1.6	32.0.203.280	Oct 7th, 2028
1.5	32.0.203.280	July 1st, 2028
1.4.1	32.0.203.259	May 13th, 2028
1.4	32.0.203.257	March 25th, 2028
1.3.1	32.0.203.242	January 17th, 2028
1.3	32.0.203.237	November 26th, 2027
1.2	32.0.201.204	July 30th, 2027

The application must verify that the user has installed NPU drivers compatible with the version of the Vitis AI EP in use.

3.7.2 APU Types

The Ryzen AI Software supports various types of NPU-enabled APUs, referred to as PHX, HPT, STX, and KRK. To programmatically determine the type of the local APU, you can enumerate the PCI devices and look for an instance with a matching Hardware ID.

Vendor	Device	Revision	APU Type
0x1022	0x1502	0x00	PHX or HPT
0x1022	0x17F0	0x00	STX
0x1022	0x17F0	0x10	STX
0x1022	0x17F0	0x11	STX
0x1022	0x17F0	0x20	KRK

The application must verify that it runs on an AMD processor with NPU, and that the Vitis AI EP version supports the NPU type.

3.7.3 NPU Utilities

When deploying applications across various NPU devices, users can determine the specific type of NPU device using Python/C++ code. Based on the detected device—such as PHX, STX, KRK, or other device—users should configure the appropriate provider options as outlined in *Model Compilation and Deployment*.

For Python, the user can get the specific NPU type using the following example *get_npu_info* function in the `%RYZEN_AI_INSTALLATION_PATH%\quicktest\quicktest.py`

For C++, a set of APIs are provided to extract information about the NPU and check driver compatibility of the VitisAI EP with the rest of the environment. For details refer to C++ NPU Utilities

3.7.4 Application Development Requirements

ONNX-RT Session

The application should only use the Vitis AI Execution Provider if the user meets the following

- The application is running on an AMD processor with an NPU type supported by the version of the Vitis AI EP in use. See *list*.
- The user installs NPU drivers compatible with the version of the Vitis AI EP they are using. See *compatibility table* .

NOTE: Sample C++ code that implements the compatibility checks to be performed before using the Vitis AI EP is available here

VitisAI EP Provider Options

For INT8 models, the application should detect the type of APU present (PHX, HPT, STX, or KRK) and set the `target` and `xclbin` provider options accordingly. Refer to the section on *using INT8 models* for more details.

For BF16 models, the application should set the `config_file` provider option to the same file that the compiler used to precompile the BF16 model. Refer to the section on *using BF16 models* for more details.

Pre-Compiled Models

Precompiled models load instantly and execute immediately on the NPU, significantly improving session creation time and overall end-user experience.

AMD recommends using the ONNXRuntime *EP Context Cache* feature for saving and reloading compiled models.

BF16 models

The deployment version of the VitisAI Execution Provider (EP) does not support the on-the-fly compilation of BF16 models. Applications utilizing BF16 models must include pre-compiled versions of these models. The VitisAI EP can then load the pre-compiled models and deploy them efficiently on the NPU.

INT8 models

The compiler recommends including pre-compiled versions of INT8 models, but it is not mandatory.

3.7.5 Application Packaging Requirements

Ryzen AI 1.8.0 beta version is available for download at the following link. This beta version requires updated NPU driver.

A C++ application built on the Ryzen AI ONNX Runtime must include the following components in its distribution package:

For INT8 models

- DLLs:
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\aiecompiler_client.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\DirectML.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\dyn_dispatch_core.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_providers_shared.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_providers_vitisai.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_vitis_ai_custom_ops.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_vitisai_ep.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime.dll
- NPU Binary files (.xclbin) from the `%RYZEN_AI_INSTALLATION_PATH%\voe-4.0-win_amd64\xclbins` folder
- Recommended but not mandatory: pre-compiled models in the form of *Onnx Runtime EP context models*

For BF16 models

- DLLs:
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_providers_shared.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_providers_vitisai.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime_vitisai_ep.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\dyn_dispatch_core.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\aiecompiler_client.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\vaiml.dll
- Pre-compiled models in the form of *Vitis AI EP cache folders*

For LLMs

- DLLs:
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime-genai.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnxruntime.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\ryzen_mm.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\onnx_custom_ops.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\libutf8_validity.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\abseil_dll.dll
  - %RYZEN_AI_INSTALLATION_PATH%\deployment\DirectML.dll

3.8 Whisper.cpp support

Ryzen AI provides NPU acceleration for whisper.cpp through an AMD-maintained fork. On AMD Ryzen AI 300 Series systems, Whisper can fully offload the encoder to the NPU, which can deliver a significant speedup versus CPU-only runs. NPU acceleration is currently supported on Windows only, with Linux support planned.

For setup steps, build flags, and NPU-optimized model guidance, refer to the GitHub README at https://github.com/amd/whisper.cpp?tab=readme-ov-file#amd-ryzen-ai-support-for-npu

3.9 LLM Deployment Overview

Large Language Models (LLMs) can be deployed on Ryzen AI PCs with NPU and GPU acceleration. NPU-only and hybrid execution modes, which utilize both the NPU and integrated GPU (iGPU), are supported by ONNXRuntime GenAI (OGA). GPU-only acceleration is enabled through llama.cpp. See the *LLM Execution Mode Comparison* below for detailed information.

3.9.1 Execution Modes

Table 5: LLM Execution Mode Comparison

Mode	Framework(s)	Compute Allocation	Primary Use Case
NPU-Only	OnnxRuntime GenAI (OGA)	Neural Processing Unit (NPU) exclusive	Maximum NPU utilization while preserving iGPU for parallel workloads
Hybrid	OnnxRuntime GenAI (OGA)	Dynamic NPU + iGPU partitioning	Interactive inference with optimal prefill/decode performance
GPU	llama.cpp	Dedicated GPU execution	High-throughput inference on discrete/integrated GPU
CPU	OGA or llama.cpp	Traditional CPU-based inference	Baseline compatibility across all processor generations

3.9.2 Hardware Requirements

Table 6: Supported Processor Configurations

Processor Series	NPU-Only	Hybrid	GPU/CPU
Ryzen AI 300 (STX/KRK)	✓	✓	✓
Ryzen AI 7000/8000			✓

Development Interfaces

The Ryzen AI LLM software stack is available through three development interfaces, each suited for specific use cases as outlined in the sections below. All three interfaces are built on top of native OnnxRuntime GenAI (OGA) libraries or llama.cpp libraries, as shown in the *Ryzen AI Software Stack* diagram below.

The high-level Python APIs, as well as the Server Interface, also leverage the Lemonade SDK,

which is multi-vendor open-source software that provides everything necessary for quickly getting started with LLMs on OGA or llama.cpp.

A key benefit of Lemonade is that software developed against their interfaces is portable to many other execution backends.

Table 7: Ryzen AI Software Stack

Your Python Application	Your LLM Stack	Your Native Application
*Lemonade Python API**	*Lemonade Server Interface**	OGA C++ Headers **OR** llama.cpp C++ Headers
Custom AMD OnnxRuntime GenAI (OGA) **OR** llama.cpp*		
AMD Ryzen AI Driver and Hardware		

* indicates open-source software (OSS).

3.9.3 Server Interface (REST API)

The Server Interface provides a convenient means to integrate with applications that:

- Already support an LLM server interface, such as the Ollama server or OpenAI API.
- Are written in any language (C++, C#, Javascript, etc.) that supports REST APIs.
- Benefits from process isolation for the LLM backend.

Lemonade Server is available in two ways:

- **Standalone Windows GUI installer**: Quick setup with a desktop shortcut for immediate use. (Recommended for end users, see *Server Interface (REST API)*)
- **Full Lemonade SDK**: Complete development toolkit with server interface included. (Recommended for developers, see *High-Level Python SDK* for Python SDK)

For example applications that have been tested with Lemonade Server, see the Lemonade Server Examples.

3.9.4 High-Level Python SDK

The high-level Python SDK, Lemonade, allows you to get started using PyPI installation in approximately 5 minutes.

This SDK allows you to:

- Experiment with models in hybrid or NPU-only execution mode on Ryzen AI hardware.
- Validate inference speed and task performance.
- Integrate with Python apps using a high-level API.

To get started in Python, follow these instructions: *High-Level Python SDK*.

3.9.5 OGA APIs for C++ Libraries and Python

Native C++ libraries for OGA are available to give full customizability for deployment into native applications. The Python bindings for OGA also provide a customizable interface for Python development.

To get started with the OGA APIs, follow these instructions: *OnnxRuntime GenAI (OGA) Flow*.

Supported LLMs

AMD provides a set of pre-optimized LLMs ready to be deployed with Ryzen AI Software and the supporting runtime for hybrid and/or NPU-only execution. These include popular architectures such as Llama-2, Llama-3, Mistral, DeepSeek Distill models, Qwen-2, Qwen-2.5, Qwen-3, Gemma-2, Phi-3, Phi-3.5, and Phi-4.

For the collection of supported models, visit *Pre-optimized Models*

It is also possible to run fine-tuned versions of the models listed (for example, fine-tuned versions of Llama2 or Llama3). For instructions on how to prepare a fine-tuned OGA model, refer to *Preparing OGA Models*.

End to End OGA Validation

A Jupyter Notebook example is provided to demonstrate end-to-end validation of OGA hybrid and NPU-only execution. This notebook includes:

- Installation
- Command Syntax
- Benchmarking
- Subjective Evaluation
- Objective Evaluation

To run the notebook, visit the Lemonade Tools Tutorial.

3.10 Server Interface (REST API)

The Lemonade SDK offers a server interface that allows your application to load an LLM on Ryzen AI hardware in a process, and then communicate with this process using standard REST APIs. This allows applications written in any language (C#, JavaScript, Python, C++, etc.) to easily integrate with Ryzen AI LLMs.

Server interfaces are used across the LLM ecosystem because they allow for no-code plug-and-play between the higher level of the application stack (GUIs, agents, RAG, etc.) with the LLM

and hardware that have been abstracted by the server. For more information, see the Understanding local LLM Servers Guide.

For example, open source projects such as *Open WebUI* have out-of-box support for connecting to a variety of server interfaces, which in turn allows users to quickly start working with LLMs in a GUI.

3.10.1 Server Setup

Lemonade Server can be installed via the Lemonade Server Installer executable by following these steps:

1. Make sure your system has the recommended Ryzen AI driver installed as described in *Install NPU Drivers*.
2. Download and install `Lemonade_Server_Installer.exe` from the latest Lemonade release.
3. Launch the server by double-clicking the `lemonade_server` shortcut added to your desktop.

For a visual walkthrough of this process, watch our Lemonade Introductory Video:

See the Lemonade Server Documentation for more details.

3.10.2 Server Usage

The Lemonade Server provides the following OpenAI-compatible endpoints:

- POST `/api/v1/chat/completions` - Chat Completions (messages to completions)
- POST `/api/v1/completions` - Text Completions (prompt to completion)
- POST `/api/v1/responses` - Chat Completions (prompt|messages -> event)
- GET `/api/v1/models` - List available models

Please refer to the server specification document for details about the request and response formats for each endpoint.

The OpenAI API documentation also has code examples for integrating streaming completions into an application.

Supported Applications

The Lemonade Server supports a variety of applications that can connect to it using the OpenAI API. Some of the applications that have been tested with Lemonade Server can be found at Lemonade Server Apps.

A short list of applications that have been tested with Lemonade Server includes:

3.10.3 Next Steps

- See Lemonade Server Examples to find applications that have been tested with Lemonade Server.
- Check out the Lemonade Server specification to learn more about supported features.
- Try out your Lemonade Server install with any application that uses the OpenAI chat completions API.

3.11 High-Level Python SDK

A Python environment offers flexibility for experimenting with LLMs, profiling them, and integrating them into Python applications. We use the Lemonade SDK to get up and running quickly.

To get started, follow these instructions.

3.11.1 System-level pre-requisites

You only need to do this once per computer:

1. Make sure your system has the recommended Ryzen AI driver installed as described in *Install NPU Drivers*.
2. Download and install Miniconda for Windows or Miniforge for Windows.
3. Launch a terminal and call `conda init`.

3.11.2 Environment Setup

To create and set up an environment, run these commands in your terminal:

```
conda create -n ryzenai-llm python=3.12
conda activate ryzenai-llm
pip install lemonade-sdk[dev,oga-ryzenai] --extra-index-url=https://pypi.
↪amd.com/simple
```

3.11.3 Validation Tools

Now that you have completed installation, you can try prompting an LLM like this (where `PROMPT` is any prompt you like).

Run this command in a terminal that has your environment activated:

```
lemonade -i amd/Llama-3.2-1B-Instruct-awq-g128-int4-asym-fp16-onnx-hybrid␣
↪oga-load --device hybrid --dtype int4 llm-prompt --max-new-tokens 64 -p␣
↪PROMPT
```

For an end-to-end example demonstrating the Validation Tools, visit the Lemonade Tools Tutorial.

3.11.4 Python API

You can also run this code to try out the high-level Lemonade API in a Python script:

```
from lemonade.api import from_pretrained

model, tokenizer = from_pretrained(
    "amd/Llama-3.2-1B-Instruct-awq-g128-int4-asym-fp16-onnx-hybrid",␣
↪recipe="oga-hybrid"
)

input_ids = tokenizer("This is my prompt", return_tensors="pt").input_ids
response = model.generate(input_ids, max_new_tokens=30)

print(tokenizer.decode(response[0]))
```

3.11.5 Next Steps

From here, you can check out the Jupyter Notebook that provides an end-to-end validation of OGA hybrid and NPU-only execution. To run the notebook, visit the Lemonade Tools Tutorial.

3.12 OnnxRuntime GenAI (OGA) Flow

Ryzen AI Software supports deploying LLMs on Ryzen AI PCs using the native ONNX Runtime Generate (OGA) C++ or Python API. The OGA API is the lowest-level API available for building LLM applications on a Ryzen AI PC. It supports the following execution modes:

- Hybrid execution mode: This mode uses both the NPU and iGPU to achieve the best TTFT and TPS during the prefill and decode phases.
- NPU-only execution mode: This mode uses the NPU exclusively for both the prefill and decode phases. Two types of NPU models are available - Token Fusion (long context) and Full Fusion (best performance). See *NPU Models: Token Fusion vs Full Fusion* for the full comparison.

3.12.1 Supported Configurations

The Ryzen AI OGA flow supports Strix and Krackan Point processors. Phoenix (PHX) and Hawk (HPT) processors are not supported.

3.12.2 Requirements

- Install NPU Drivers and Ryzen AI MSI installer. See *Installation Instructions* for more details.
- Install GPU device driver: Ensure GPU device driver https://www.amd.com/en/support is installed
- Install Git for Windows (needed to download models from HF): https://git-scm.com/downloads

3.12.3 Pre-optimized Models

AMD provides a set of pre-optimized LLMs ready to be deployed with Ryzen AI Software and the supporting runtime for hybrid and/or NPU-only execution. These include popular architectures such as Llama-2, Llama-3, Mistral, DeepSeek Distill models, Qwen-2, Qwen-2.5, Qwen-3, Gemma-2, Gemma-3, GPT-OSS, Phi-3, Phi-3.5, and Phi-4.

For the complete list of supported pre-optimized models and their available variants, see llm_list.

- Hugging Face collection of **hybrid models**: https://huggingface.co/collections/amd/ryzen-ai-180-hybrid
- Hugging Face collection of **NPU long-context models** (up to 16K, Token Fusion): https://huggingface.co/collections/amd/ryzen-ai-180-npu-16k
- Hugging Face collection of **NPU best-performance models** (up to 4K, Full Fusion): https://huggingface.co/collections/amd/ryzen-ai-180-npu-4k

NPU Models: Token Fusion vs Full Fusion

AMD provides two types of NPU models. Choose based on your use case: Token Fusion for long-context workloads, or Full Fusion for higher throughput on shorter sequences.

	Token Fusion	Full Fusion
Max context (input + output)	Up to 16K tokens	Up to 4096 tokens
Best for	Long-context workloads	Higher throughput on shorter sequences

Each OGA model folder contains a `genai_config.json` file, which holds configuration settings for the model. The `session_option` section is where specific runtime dependencies are specified.

3.12.4 Changes Compared to Previous Release

- OGA version is updated to v0.14.0 (Ryzen AI 1.8) from v0.11.2 (Ryzen AI 1.7.1).
- For the 1.8 release, a new set of hybrid and NPU models has been published. Models from earlier releases are not compatible with this version. Download the updated models.

3.12.5 Compatible OGA APIs

Pre-optimized hybrid or NPU LLMs can be executed using the official OGA C++ and Python APIs. The current release is compatible with OGA version 0.14.0. For detailed documentation and examples, refer to the official OGA repository: https://github.com/microsoft/onnxruntime-genai/tree/rel-0.14.0

3.12.6 LLMs Test Programs

The Ryzen AI installation includes test programs (in C++ and Python) that can be used to run LLMs and understand how to integrate them in your application.

The steps for deploying the pre-optimized models using the sample programs are described in the following sections.

Steps to run C++ program and sample python script.

1. (Optional) Enable Performance Mode

To run LLMs in best performance mode, follow these steps:

- Go to `Windows` → `Settings` → `System` → `Power`, and set the power mode to **Best Performance**.
- Open a terminal and run:

```
cd C:\Windows\System32\AMD
xrt-smi configure --pmode performance
```

2. Activate the Ryzen AI Conda Environment and install `torch` library.

Run the following commands:

```
conda activate ryzen-ai-<version>
```

This step is required for running the python script.

> **Note**
>
> For the C++ program, if you choose not to activate the Conda environment, open a Windows Command Prompt and manually set the environment variable before continuing:

```
set RYZEN_AI_INSTALLATION_PATH=C:\\Program Files\\RyzenAI\\<version>
```

C++ Program

Use the `model_benchmark.exe` executable to test LLMs and identify DLL dependencies for C++ applications.

> **Note**
>
> `model_benchmark.exe` is for **performance** measurement only. It feeds the model a raw or synthetic prompt without applying a chat template, so its generated text is often repetitive and is **not** meant to reflect output quality. For **instruct/chat models**, use the chat-template scripts to evaluate accuracy or produce coherent responses (see *Python Script (with Chat Template)*).

1. Set Up a working directory and copy required Files

```
mkdir llm_run
cd llm_run

:: Copy the sample C++ executable
xcopy /Y "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\model_benchmark.exe" .

:: Copy the sample prompt file
xcopy /Y "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\amd_genai_prompt.txt" .

:: Copy required DLLs
xcopy /Y "%RYZEN_AI_INSTALLATION_PATH%\deployment\." .
```

2. Download model from Hugging Face

```
:: Install Git LFS if you haven't already: https://git-lfs.com
git lfs install

:: Clone the model repository
git clone https://huggingface.co/amd/Llama-2-7b-chat-hf-onnx-ryzenai-
↪hybrid
```

3. Run `model_benchmark.exe`

Provide the prompt using **either** `--prompt_file` (a prompt text file) **or** `-l` (a synthetic prompt of the given token length). These options are mutually exclusive.

```
:: Using a synthetic prompt length:
.\model_benchmark.exe -i <path_to_model_dir> -l <prompt_length> -g
```

(continues on next page)

(continued from previous page)

```
↪<generation_length>
:: Example:
.\model_benchmark.exe -i Llama-2-7b-chat-hf-onnx-ryzenai-hybrid -l 1024 -
↪g 128

:: Using a prompt file:
.\model_benchmark.exe -i <path_to_model_dir> --prompt_file <prompt_file> -
↪g <generation_length>
:: Example:
.\model_benchmark.exe -i Llama-2-7b-chat-hf-onnx-ryzenai-hybrid --prompt_
↪file amd_genai_prompt.txt -g 128
```

Key options: Run `model_benchmark.exe --help` for the complete list of options.

Long Context Support

Ryzen AI supports long context (beyond 4096 tokens) for **Hybrid models** and **Token Fusion NPU models**.

Token Fusion NPU Models

Token Fusion NPU models are pre-built with long context support up to 16K tokens. No additional configuration is required — simply download the model from Hugging Face and run it.

```
:: Example: Clone a Token Fusion NPU model
git clone https://huggingface.co/amd/Phi-3.5-mini-instruct-onnx-ryzenai-
↪npu

:: Run with long context (synthetic 16000-token prompt)
.\model_benchmark.exe -i <path_to_model_dir> -l 16000 -g 128
```

Hybrid Models

If the total number of tokens exceeds 4096 for a hybrid model, follow the steps below.

Steps to run long context:

1. Make the following changes in `genai_config.json` file.
   - Add `"hybrid_opt_chunk_context": "1"` under `model.decoder.session_options.provider_options.RyzenAI`.

```
{
"model": {
   "bos_token_id": 1,
```

(continues on next page)

(continued from previous page)

```
  "context_length": 16384,
  "decoder": {
        "session_options": {
        "log_id": "onnxruntime-genai",
        "provider_options": [
        {
           "RyzenAI": {
              "external_data_file": "model_jit.pb.bin",
              "hybrid_opt_free_after_prefill": "1",
              "hybrid_opt_max_seq_length": "4096",
                                          "hybrid_opt_chunk_
↪context": "1"
           }
        }
        ]
     },
```

- Add `"chunk_size":2048` under `search`.

```
"search": {
     "diversity_penalty": 0.0,
     "do_sample": false,
     "chunk_size": 2048,
     ...
```

2. Run the model using `model_benchmark.exe` with a synthetic long-context prompt.

```
:: Generate a 16000-token prompt and 128 output tokens
.\model_benchmark.exe -i <path_to_model_dir> -l 16000 -g 128
```

Note

The sample test application `model_benchmark.exe` accepts `-l` for input token length and `-g` for output token length.

- **Full Fusion NPU models** support up to 4096 tokens in total (input + output). By default, `-g` is set to 128. If the input length is close to 4096, you must adjust `-g` so the sum of input and output tokens does not exceed 4096. For example, `-l 4000 -g 96` is valid (4000 + 96 4096), while `-l 4000 -g 128` will exceed the limit and result in an error.
- **Token Fusion NPU models** support long context up to 16K tokens (input + output) with no additional configuration.
- **Hybrid models**: The combined number of input and output tokens must not exceed the

model's `context_length`. You can verify the `context_length` in the `genai_config.json` file. For example, if a model's `context_length` is 8,000, the total token count (input + output) must not exceed 8,000.

The long context feature has been tested for Token Fusion NPU models and Hybrid models up to 16,000 tokens.

Python Script

This section uses a basic `run_model.py` sample script.

Note

`run_model.py` does not apply a chat template. If your model uses a chat template, use *Python Script (with Chat Template)* (`model_chat.py`) instead for more accurate output.

1. Navigate to your working directory and download model.

```
:: Install Git LFS if you haven't already: https://git-lfs.com
git lfs install

:: Clone the model repository
git clone https://huggingface.co/amd/Llama-2-7b-chat-hf-onnx-ryzenai-
↪hybrid
```

2. Run sample python script

```
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\run_model.py" -m <model_
↪folder> -l <max_length>

:: Example command
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\run_model.py" -m "Llama-
↪2-7b-chat-hf-onnx-ryzenai-hybrid" -l 256
```

Python Script (with Chat Template)

For models that use chat templates, the `model_chat.py` script provides better output quality by automatically loading and applying the chat template from the model folder during inference. The script also supports single-prompt, multi-turn context cache testing, and interactive chat with timing output.

The script is included in the Ryzen AI installation:

```
:: Single prompt with timing
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\model_chat.py" -m <model_
↪folder> -pr amd_genai_prompt.txt --timings

:: Long context support (increase context window to e.g. 16k)
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\model_chat.py" -m <model_
↪folder> -pr amd_genai_prompt_long.txt -mpt 16000

:: Interactive chat
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\model_chat.py" -m <model_
↪folder>
```

For the full list of options including multi-turn JSON testing, guided generation, and advanced flags, refer to the RyzenAI-SW repository.

It is highly recommended to use `model_chat.py` for the GPT-OSS-20B NPU model.

3.12.7 Vision Language Model (VLM)

AMD provides a pre-optimized Gemma-3-4b-it multimodal model ready to be deployed with Ryzen AI Software. Support for this model is available starting with the Ryzen AI 1.7 release.

Model: Gemma-3-4b-it

VLM inference requires dedicated Python scripts, which are included in the Ryzen AI installation at `%RYZEN_AI_INSTALLATION_PATH%\LLM\example\vlm`.

Quick Inference

Use `vlm_run.py` to quickly test a model and see output:

```
:: Basic inference
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\vlm\vlm_run.py" -m
↪<model_folder> -i <image_path>

:: Custom prompt
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\vlm\vlm_run.py" -m
↪<model_folder> -i <image_path> -p "What's in this image?"

:: Resize image before running
python "%RYZEN_AI_INSTALLATION_PATH%\LLM\example\vlm\vlm_run.py" -m
↪<model_folder> -i <image_path> --image_size 1024 1024
```

For benchmarking scripts (`vlm_benchmark.py`, `run_all_benchmarks.py`) and detailed options, refer to the README in the `vlm` directory or the RyzenAI-SW repository.

3.12.8 Building C++ Applications

The RyzenAI-SW repository provides a complete C++ example demonstrating how to load a pre-optimized OGA model, apply the chat template, and run inference using the OGA C++ API. It includes the full source and CMake build instructions.

Repository: https://github.com/amd/RyzenAI-SW/tree/main/LLM-examples/oga_api

3.12.9 Using Fine-Tuned Models

It is also possible to run fine-tuned versions of the pre-optimized OGA models.

To do this, the fine-tuned models must first be prepared for execution with the OGA flow. For instructions on how to do this, refer to the page about *Preparing OGA Models*.

After a fine-tuned model has been prepared for execution, it can be deployed by following the steps described previously in this page.

3.12.10 Running LLM with pip install

In addition to the full RyzenAI software stack, we also provide standalone wheel files for the users who prefer using their own environment. To prepare an environment for running the Hybrid and NPU-only LLM independently, perform the following steps:

1. Create a new python environment and activate it.

```
conda create -n <env_name> python=3.12 -y
conda activate <env_name>
```

2. Install onnxruntime-genai wheel file.

```
pip install --extra-index-url https://pypi.amd.com/ryzenai_llm/1.8.0/
↪windows/simple onnxruntime-genai-directml-ryzenai==0.14.0
pip install --extra-index-url https://pypi.amd.com/ryzenai_llm/1.8.0/
↪windows/simple model-generate==1.8.0
```

3. Navigate to your working directory and download the desired Hybrid/NPU model

```
cd working_directory
git clone <link_to_model>
```

4. Run the Hybrid or NPU model.

3.13 Preparing OGA Models

This section describes the process for preparing LLMs for deployment on a Ryzen AI PC using the hybrid or NPU-only execution mode. Currently, the flow supports only fine-tuned versions of the models already supported (as listed in *OnnxRuntime GenAI (OGA) Flow* page). For example, fine-tuned versions of Llama2 or Llama3 can be used. However, different model families with architectures not supported by the hybrid flow cannot be used.

For fine-tuned models that introduce architectural changes requiring new operator shapes not available in the Ryzen AI runtime, refer to the oga_op_prepare

Preparing a LLM for deployment on a Ryzen AI PC involves 2 steps:

1. **Quantization**: The pretrained model is quantized to reduce memory footprint and better map to compute resources in the hardware accelerators
2. **Postprocessing**: During the postprocessing the model is exported to OGA followed by NPU-only or Hybrid execution mode specific postprocess to obtain the final deployable model.

3.13.1 Quantization

Prerequisites

A Linux machine with AMD (e.g., AMD Instinct MI Series) or Nvidia GPUs

Setup

1. Create and activate Conda Environment

```
conda create --name <conda_env_name> python=3.12
conda activate <conda_env_name>
```

2. If Using AMD GPUs, update PyTorch to use ROCm

```
pip3 install torch torchvision torchaudio --index-url https://download.
↪pytorch.org/whl/rocm6.1
python -c "import torch; print(torch.cuda.is_available())" # Must return␣
↪`True`
```

3. Download `AMD Quark 0.11` and unzip the archive
4. Install Quark:

```
cd <extracted amd quark-version>
pip install amd_quark-<version>+<>.whl
```

5. Install other dependencies

```
pip install datasets
pip install transformers==4.57.6
pip install accelerate
pip install evaluate
pip install nltk
```

Generate Quantized Model

Use following command to run Quantization. In a GPU equipped Linux machine the quantization can take about 30-60 minutes.

```
cd examples/torch/language_modeling/llm_ptq/

python quantize_quark.py \
     --no_trust_remote_code \
     --model_dir "meta-llama/Llama-2-7b-chat-hf"  \
     --output_dir <quantized safetensor output dir>  \
     --quant_scheme uint4_wo_128 \
     --num_calib_data 128 \
     --seq_len 512 \
     --quant_algo awq \
     --dataset pileval_for_awq_benchmark \
     --model_export hf_format \
     --data_type <datatype> \
     --exclude_layers []
```

- Use `--data_type bfloat16` for bf16 pretrained model. For fp32/fp16 pretrained model use `--datatype float16`
- Quark natively supports AWQ quantization for popular architectures. If AWQ is not supported by default, you must create an AWQ configuration file and pass it to the `quantize_quark.py` script using the `--quant_algo_config_file awq <custom awq config json>` option. For details on creating an AWQ config file, see the Quark documentation AWQ example: https://quark.docs.amd.com/latest/tutorials/torch/example_awq.html
- Not using `--exclude_layers` parameter may result in model-specific defaults which may exclude certain layers like output layers.
- To specify a group size other than 128, such as 32, use `--quant_scheme uint4_wo_32` instead of `--quant_scheme uint4_wo_128`. Available group sizes are 32, 64, and 128 (e.g., uint4_wo_32, uint4_wo_64, uint4_wo_128)
- Quark supports quantizing layers with different group sizes, use `--layer_quant_scheme lm_head uint4_wo_32` to quantize the model with 32 group size for lm_head

The quantized model is generated in the <quantized safetensor output dir> folder.

Note: For the Phi-4 model, the following quantization recipe is recommended for better accuracy:

- Use `--quant_algo gptq`
- Add `--layer_quant_scheme lm_head uint4_wo_32`

Note:: Currently the following files are not copied into the quantized model folder and must be copied manually:

- For Phi-4 models: `configuration_phi3.py`
- For ChatGLM-6b models: `tokenizer.json`

3.13.2 Postprocessing

Copy the quantized model to the Windows PC with Ryzen AI installed, and activate the Ryzen AI Conda environment.

```
conda activate ryzen-ai-<version>
```

Hybrid Execution Mode

Generate the final model for Hybrid execution mode (NPU prefill phase + GPU token phase):

```
model_generate --hybrid --input <quantized_model_path> --output <output_
↪dir>
```

NPU Execution Mode

Several NPU optimization levels are available depending on model support and performance requirements. To check which variants (full fusion, token fusion, hybrid) are supported for a given model architecture, see llm_list.

Full Fusion (**Best performance**, recommended for supported models):

```
model_generate --npu --full_fusion --input <quantized_model_path> --
↪output <output_dir>
```

Token Fusion (**long-context** supportted):

```
:: 16K context
model_generate --npu --token_fusion --input <quantized_model_path> --
↪output <output_dir> --extra_options max_seq_len=16384
```

Basic (safe default for new or untested models):

```
model_generate --npu --basic --input <quantized_model_path> --output
↪<output_dir>
```

Eager:

```
model_generate --npu --eager --input <quantized_model_path> --output
↪<output_dir>
```

OGA Export Only

To export the quantized model to OGA format without performing any NPU or Hybrid postprocessing:

```
model_generate --oga_only --input <quantized_model_path> --output <output_
↪dir>
```

Memory Optimization

Add `--mem_optimize` to any recipe to optimize for 16 GB laptop configurations:

```
model_generate --hybrid --mem_optimize --input <quantized_model_path> --
↪output <output_dir>
```

```
model_generate --npu --token_fusion --mem_optimize --input <quantized_
↪model_path> --output <output_dir>
```

Note: During the `model_generate` step, the quantized model is first converted to an OGA model using ONNX Runtime GenAI Model Builder (version 0.14.0). It is possible to use a standalone environment for exporting an OGA model, refer to the official ONNX Runtime GenAI Model Builder documentation. Once you have an exported OGA model, you can pass it directly to the `model_generate` command with `--input`, which will skip the export step and perform only the post-processing.

Here are simple commands to export an OGA model from a quantized model using a standalone environment:

```
conda create --name oga_builder_env python=3.12
conda activate oga_builder_env


pip install onnxruntime-genai==0.14.0
# pip install other necessary packages
pip install ....


python3 -m onnxruntime_genai.models.builder -m <input quantized model> -o
↪<output OGA model> -p int4 -e dml
```

3.14 DirectML Flow

3.14.1 Prerequisites

- DirectX12 capable Windows OS (Windows 11 recommended)
- Latest AMD GPU device driver installed
- Microsoft Olive for model conversion and optimization
- Latest ONNX Runtime DirectML EP

You can ensure GPU driver and DirectX version from `Windows Task Manager` -> `Performance` -> `GPU`

3.14.2 Running models on Ryzen AI GPU

Running models on the Ryzen AI GPU is accomplished in two simple steps:

Model Conversion and Optimization: After the model is trained, Microsoft Olive Optimizer can be used to convert the model to ONNX and optimize it for optimal target execution.

For additional information, refer to the Microsoft Olive Documentation

Deployment: Once the model is in the ONNX format, the ONNX Runtime DirectML EP (`DmlExecutionProvider`) is used to run the model on the AMD Ryzen AI GPU.

For additional information, refer to the ONNX Runtime documentation for the DirectML Execution Provider

3.14.3 Examples

- Optimizing and running ResNet on Ryzen AI GPU

3.14.4 Additional Resources

- Article on how AMD and Black Magic Design worked together to accelerate Davinci Resolve Studio workload on AMD hardware:
  - AI Accelerated Video Editing with DaVinci Resolve 18.6 & AMD Radeon Graphics
- Blog posts on using the Ryzen AI Software for various generative AI workloads on GPU:
  - Automatic1111 Stable Diffusion WebUI with DirectML Extension on AMD GPUs
  - Running Optimized Llama2 with Microsoft DirectML on AMD Radeon Graphics

- AI-Assisted Mobile Workstation Workflows Powered by AMD Ryzen™ AI

3.15 NPU Management Interface

3.15.1 Introduction

The `xrt-smi` utility is a command-line interface to monitor and manage the NPU integrated AMD CPUs.

On the Windows platform, `xrt-smi` is installed in `C:\Windows\System32\AMD`. Ensure that the following path `C:\Windows\System32\AMD` is set in the System PATH variable. This will allow it to be directly invoked from within the conda environment created by the Ryzen AI Software installer.

The `xrt-smi` utility currently supports three primary commands:

- `examine` - generates reports related to the state of the AI PC and the NPU.
- `validate` - executes sanity tests on the NPU.
- `configure` - manages the performance level of the NPU.

By default, the output of the `xrt-smi examine` and `xrt-smi validate` commands goes to the terminal. It can also be written to file in JSON format as shown below:

```
xrt-smi examine -f JSON -o <path/to/output.json>
```

The utility also support the following options which can be used with any command:

- `--help` - help to use xrt-smi or one of its sub commands
- `--version` - report the version of XRT, driver and firmware
- `--verbose` - turn on verbosity
- `--batch` - enable batch mode (disables escape characters)
- `--force` - when possible, force an operation. Eg - overwrite a file in examine or validate

In Windows, the `xrt-smi` utility requires Microsoft Visual C++ Redistributable (version 2015 to 2022) to be installed.

3.15.2 Overview of Key Commands

Command	Description
examine	system config, device name
examine –report platform	performance mode, power
examine –report aie-partitions	hw contexts
validate –run latency	latency test
validate –run throughput	throughput test
validate –run gemm	INT8 GEMM test TOPS. This is a full array test and it should not be run while another workload is running. **NOTE**: This command is not supported on PHX and HPT NPUs.
configure –pmode <mode>	set performance mode

NOTE: The `examine --report aie-partition` report runtime information. These commands should be used when a model is running on the NPU. You can run these commands in a loop to see live updates of the reported data.

3.15.3 xrt-smi examine

System Information

Reports OS/system information of the AI PC and confirm the presence of the AMD NPU.

```
xrt-smi examine
```

Sample Command Line Output:

```
System Configuration
  OS Name               : Windows NT
  Release               : 26100
  Machine               : x86_64
  CPU Cores             : 20
  Memory                : 32063 MB
  Distribution          : Microsoft Windows 11 Enterprise
  Model                 : HP OmniBook Ultra Laptop 14-fd0xxx
  BIOS Vendor           : HP
  BIOS Version          : W81 Ver. 01.01.14

XRT
  Version               : 2.19.0
  Branch                : HEAD
  Hash                  : f62307ddadf65b54acbed420a9f0edc415fefafc
  Hash Date             : 2025-03-12 16:34:48
```

(continues on next page)

(continued from previous page)

```
  NPU Driver Version   : 32.0.203.257
  NPU Firmware Version : 1.0.7.97

Device(s) Present
|BDF              |Name       |
|-----------------|-----------|
|[00c4:00:01.1]   |NPU Strix  |
```

Sample Command Line Output in Linux:

```
> xrt-smi examine
  System Configuration
      OS Name              : Linux
      Release              : 6.11.0-26-generic
      Machine              : x86_64
      CPU Cores            : 24
      Memory               : 31440 MB
      Distribution         : Ubuntu 24.04 LTS
      GLIBC                : 2.39
      Model                : BIRMANPLUS
      BIOS Vendor          : AMD
      BIOS Version         : TXB1001dB

  XRT
      Version              : 2.20.0
      Branch               : master
      Hash                 : 7a277facefecab6c87ac835916021c63d2e395dd
      Hash Date            : 2025-06-16 21:28:35
      amdxdna              : 2.20.0_20250617,␣
↪e7233301f8e4d8d1b1678f3dc3492c826290e314
      NPU Firmware Version : 255.0.1.5

  Device(s) Present
  |BDF              |Name       |
  |-----------------|-----------|
  |[0000:c5:00.1]   |NPU Strix  |
```

Sample JSON Output:

```
{
"schema_version": {
    "schema": "JSON",
    "creation_date": "Tue Mar 18 22:43:38 2025 GMT"
```

(continues on next page)

(continued from previous page)


```
},
"system": {
    "host": {
        "os": {
            "sysname": "Windows NT",
            "release": "26100",
            "machine": "x86_64",
            "distribution": "Microsoft Windows 11 Enterprise",
            "model": "HP OmniBook Ultra Laptop 14-fd0xxx",
            "hostname": "XCOUDAYD02",
            "memory_bytes": "0x7d3f62000",
            "cores": "20",
            "bios_vendor": "HP",
            "bios_version": "W81 Ver. 01.01.14"
        },
        "xrt": {
            "version": "2.19.0",
            "branch": "HEAD",
            "hash": "f62307ddadf65b54acbed420a9f0edc415fefafc",
            "build_date": "2025-03-12 16:34:48",
            "drivers": [
                {
                    "name": "NPU Driver",
                    "version": "32.0.203.257"
                }
            ]
        },
        "devices": [
            {
                "bdf": "00c4:00:01.1",
                "device_class": "Ryzen",
                "name": "NPU Strix",
                "id": "0x0",
                "firmware_version": "1.0.7.97",
                "instance": "mgmt(inst=1)",
                "is_ready": "true"
            }
        ]
    }
  }
}
```

Platform Information

Reports more detailed information about the NPU, such as the performance mode and power consumption.

```
xrt-smi examine --report platform
```

Sample Command Line Output:

```
-------------------------
[00c5:00:01.1] : NPU Strix
-------------------------
Platform
  Name                   : NPU Strix
  Power Mode             : Default

Estimated Power          : 1.277 Watts
```

NOTE: Power reporting is not supported on PHX and HPT NPUs. Power reporting is only available on STX devices and onwards. Report "Estimated Power" is currently unavailable for Linux users.

NPU Partitions

Reports details about the NPU partition and column occupancy on the NPU.

```
xrt-smi examine --report aie-partitions
```

Sample Command Line Output:

```
-------------------------
[00c5:00:01.1] : NPU Strix
-------------------------
AIE Partitions
  Partition Index: 0
    Columns: [0, 1, 2, 3]
    HW Contexts:
      |PID    |Ctx ID  |Status  |Instr BO  |Sub  |Compl  |Migr  |Err ␣
↪|Suspensions  |Prio     |GOPS   |EGOPS  |FPS  |Latency  |
      |-------|--------|--------|----------|-----|-------|------|-----|---
↪----------|---------|-------|-------|-----|---------|
      |20696  |0       |Active  |64 KB     |57   |56     |0     |0    |0 ␣
↪          | Normal  |0      |0      |0    |0        |
```

NPU Context Bindings

Reports details about the columns to NPU HW context binding.

```
xrt-smi examine --report aie-partitions --verbose
```

Sample Command Line Output:

```
Verbose: Enabling Verbosity
Verbose: SubCommand: examine

--------------------------
[00c5:00:01.1] : NPU Strix
--------------------------
AIE Partitions
  Partition Index: 0
    Columns: [0, 1, 2, 3]
    HW Contexts:
      |PID    |Ctx ID  |Status  |Instr BO  |Sub  |Compl  |Migr  |Err ␣
↪|Prio    |GOPS  |EGOPS  |FPS  |Latency  |
      |-------|--------|--------|----------|-----|-------|------|-----|---
↪-----|------|-------|-----|---------|
      |20696  |0       |Active  |64 KB     |57   |56     |0     |0   ␣
↪|Normal  |0     |0      |0    |0        |

AIE Columns
  |Column  ||HW Context Slot  |
  |--------||-----------------|
  |0       ||[1]              |
  |1       ||[1]              |
  |2       ||[1]              |
  |3       ||[1]              |
```

3.15.4 xrt-smi validate

Executing a Sanity Check on the NPU

Runs a set of built-in NPU sanity tests which includes latency, throughput, and gemm.

Note: All tests are run in performance mode.

- `latency` - this test executes a no-op control code and measures the end-to-end latency on all columns
- `throughput` - this test loops back the input data from DDR through a MM2S Shim DMA channel back to DDR through a S2MM Shim DMA channel. The data movement within the

AIE array follows the lowest latency path, which means the movement is restricted to just the Shim tile.

- `gemm` - An INT8 GeMM kernel is deployed on all 32 cores by the application. Each core is storing cycle count in the core data memory. The cycle count is read by the firmware. The TOPS application uses the "XBUTIL" tool to capture the IPUHCLK while the workload runs. Once all cores are executed, the cycle count from all cores will be synced back to the host. Finally, the application uses IPUHCLK, core cycle count, and GeMM kernel size to calculate the TOPS. This is a full array test and it should not be run while another workload is running. **NOTE**: This command is not supported on PHX and HPT NPUs.
- `all` - All applicable validate tests will be executed (default)

```
xrt-smi validate --run all
```

NOTE: Some sanity checks might fail if other applications (for example MEP, Microsoft Experience Package) are also using the NPU.

Sample Command Line Output:

```
Validate Device           : [00c4:00:01.1]
   Platform               : NPU Strix
   Power Mode             : Performance
-------------------------------------------------------------------------
 ↪-----
Test 1 [00c4:00:01.1]     : gemm
   Details                : TOPS: 51.3
   Test Status            : [PASSED]
-------------------------------------------------------------------------
 ↪-----
Test 2 [00c4:00:01.1]     : latency
   Details                : Average latency: 84.2 us
   Test Status            : [PASSED]
-------------------------------------------------------------------------
 ↪-----
Test 3 [00c4:00:01.1]     : throughput
   Details                : Average throughput: 59891.0 ops
   Test Status            : [PASSED]
-------------------------------------------------------------------------
 ↪-----
Validation completed. Please run the command '--verbose' option for more␣
 ↪details
```

3.15.5 xrt-smi configure

Managing the Performance Level of the NPU

To set the performance level of the NPU, you can choose from the following modes: powersaver, balanced, performance, or default. Use the command below:

```
xrt-smi configure --pmode <default | powersaver | balanced | performance␣
↪| turbo>
```

- `default` - adapts to the Windows Power Mode setting, which can be adjusted under System -> Power & battery -> Power mode. For finer control of the NPU settings, it is recommended to use the xrt-smi mode setting, which overrides the Windows Power mode and ensures optimal results.
- `powersaver` - configures the NPU to prioritize power saving, preserving laptop battery life.
- `balanced` - configures the NPU to provide a compromise between power saving and performance.
- `performance` - configures the NPU to prioritize performance, consuming more power.
- `turbo` - configures the NPU for maximum performance performance, requires AC power to be plugged in otherwise uses `performance` mode.

Example: Setting the NPU to high-performance mode

```
xrt-smi configure --pmode performance
```

To check the current power mode, use the following command:

```
xrt-smi examine --report platform
```

3.16 AI Analyzer

AMD AI Analyzer is a tool that supports analysis and visualization of model compilation and inference on Ryzen AI.

3.16.1 Installation

If you installed the Ryzen AI software using automatic installer, AI Analyzer is already installed in the conda environment.

If you manually installed the software, you need to install the AI Analyzer wheel file in your environment.

```
python -m pip install %RYZEN_AI_INSTALLATION_PATH%\aianalyzer-<version>.
↪whl
```

AI Analyzer currently supports BF16 Models only. INT8 model analysis is not supported in this version, but we plan to expand support in a future release.

3.16.2 Enabling Profiling and Visualization

The user can enable Profiling and Visualization by passing additional provider options to the ONNXRuntime Inference Session.

```
provider_options = [{
             'config_file': 'vaip_config.json',
             'cacheDir': str(cache_dir),
             'cacheKey': 'modelcachekey',
             'ai_analyzer_visualization': True,
             'ai_analyzer_profiling': True,
         }]
session = ort.InferenceSession(model.SerializeToString(),␣
↪providers=providers,
                               provider_options=provider_options)
```

The `ai_analyzer_profiling` flag enables generation of artifacts related to the inference profile. The `ai_analyzer_visualization` flag enables generation of artifacts related to graph partitions and operator fusion. The system generates these artifacts as JSON files in the current run directory.

AI Analyzer also supports native ONNX Runtime profiling, which you can use to analyze the parts of the session that run on the CPU. You can enable ONNX Runtime profiling through session options and pass it along with the provider options, as shown here:

```
 # Configure session options for profiling
 sess_options = ort.SessionOptions()
 sess_options.enable_profiling = True

 provider_options = [{
              'config_file': 'vaip_config.json',
              'cacheDir': str(cache_dir),
              'cacheKey': 'modelcachekey',
              'ai_analyzer_visualization': True,
              'ai_analyzer_profiling': True,
          }]

session = ort.InferenceSession(model.SerializeToString(), sess_options,␣
↪providers=providers,
                               provider_options=provider_options)
```

3.16.3 Launching AI Analyzer

After the compiler generates the artifacts, you can invoke *aianalyzer* through the command line as follows:

```
aianalyzer <logdir> <additional options>
```

Positional Arguments

`logdir`: Path to the folder containing generated artifacts

Additional Options

`-v`, `--version`: Show the version info and exit.

`-b ADDR`, `--bind ADDR`: Hostname or IP address on which to listen, default is 'localhost'.

`-p PORT`, `--port PORT`: TCP port on which to listen, default is '8000'.

`-n`, `--no-browser`: Prevents the opening of the default url in the browser.

`-t TOKEN`, `--token TOKEN`: Token used for authenticating first-time connections to the server. The default is to generate a new, random token. Setting to an empty string disables authentication altogether, which is not recommended.

3.16.4 Features

AI Analyzer provides visibility into how the compiler compiles and executes your AI model on Ryzen AI hardware. Its two main use cases are:

Analyzing how the compiler partitions and maps the model onto Ryzen AI's CPU and NPU accelerator 2. Profiling model performance as it executes inferencing workloads

When launched, the AI Analyzer server scans the folder specified with the logdir argument. It detects and loads all files relevant to compilation and/or

You can instruct the AI Analyzer server to either start a browser on the same host or return a URL. You can then load that URL into a browser on any host.

3.16.5 User Interface

AI Analyzer has the following three sections as seen in the left-panel navigator:

PARTITIONING - The compiler assigned a breakdown of your model to execute inference across CPU and NPU. 2. NPU INSIGHTS - A detailed look at how the compiler optimizes your model for inference execution on NPU 3. PERFORMANCE - A breakdown of inference execution through the model

The following sections describe these sections in more detail:

PARTITIONING

This section comprises two pages: Summary and Graph

Summary

The Summary page shows how the compiler assigns model operators to Ryzen's CPU and NPU, with charts of GigaOp (GOP) offloading by operator type.

There is also table titled "CPU Because" that shows the reasons why certain operators were not offloaded to the NPU.

Graph

The graph page shows an interactive diagram of the partitioned ONNX model, showing graphically how the compiler assigns the layers to the Ryzen hardware.

Toolbar

- You can choose to show/hide individual NPU partitions, if any, with the **Filter by Partition** button
- You can show or hide a panel that displays properties for selected objects through the **Show Properties** toggle button
- You can show or hide the model table through the **Show Table** toggle button.
- Settings
  - Show Processor separates operators that run on CPU and NPU respectively
  - Show Partition separates operators running on the NPU by their respective NPU partition, if any
  - Show Instance Name displays the full hierarchical name for the operators in the ONNX model

All objects in the graph have properties that you can view to the right of the graph.

Model Table

This table following the graph lists all objects in the partitioned ONNX model:

- Processor (NPU or CPU)
- Function (Layer)
- Operator
- Ports
- NPU Partitions

NPU INSIGHTS

This section comprises three pages: Summary, Original Graph, and Optimized Graph.

Summary

The Summary page gives an overview of how your model was mapped to the AMD Ryzen NPU. Charts display statistics on the number of operators and total GMACs that have been mapped to the NPU (and, if necessary, back to the CPU via the *Failsafe CPU* mechanism). The statistics are shown per operator type and NPU partition.

Original Graph

This is an interactive graph representing your model, lowered to supported NPU primitive operators and divided into partitions if necessary. As with the PARTITIONING graph, a companion table lists all model elements and supports cross-probing with the graph view. The objects in both the graph and the table also cross-probe with the PARTITIONING graph.

Toolbar

You can choose to show/hide individual NPU partitions, if any, with the **Filter by Partition** button The user can show or hide a panel that displays properties for selected objects using the **Show Properties** toggle button. The user can show or hide a code viewer showing the MLIR source code with cross-probing by clicking the **Show Code View** button. You can show and hide the following table using the **Show Table** toggle button. You can access display options for the graph with the **Settings** button.

Optimized Graph

This page shows the final model that is mapped to the NPU after all transformations and optimizations such as fusion and chaining. It also reports the operators that had to be moved back to the CPU through the *Failsafe CPU* mechanism. As usual, there is a companion table below that contains all of the graph's elements, and cross-selection is supported to and from the PARTITIONING graph and the Original Graph.

Toolbar

You can choose to show/hide individual NPU partitions, if any, with the **Filter by Partition** button The user can show or hide the panel that displays properties for selected objects using the **Show Properties** toggle button. The following table can be shown and hidden by the user using the **Show Table** toggle button. The user can access the display options for the graph with the **Settings** button.

PERFORMANCE

Use this section to view the performance of your model on RyzenAI when running one or more inferences. It comprises two pages: Summary and Timeline.

Summary

The performance summary page displays several overall statistics for the inference(s), along with charts that break down operator runtime by operator. When you enable the ONNX Runtime profiler,

you see the total inference time—including layers executed on the CPU. When you enable NPU profiling with the *ai_analyzer_profiling* flag, the tool displays additional NPU-specific statistics, including GOP and MAC efficiency, and a chart showing runtime per NPU operator type.

The clock frequency field shows the assumed NPU clock frequency, but it is editable. When you change the frequency, the system adjusts all timestamp data—collected as clock cycles but displayed in time units—accordingly.

Timeline

The Performance timeline shows a layer-by-layer breakdown of your model's execution. The upper section is a graphical depiction of layer execution across a timeline, while the lower section shows the same information in tabular format. It is important to note that the Timeline page shows one inference at a time, so if you have captured profiling data for two or more inferences, you can choose which one to display with the **Inferences** chooser.

Within each inference, you can examine the overall model execution or the detailed NPU execution data by using the **Partition** chooser.

Toolbar

The **Show Properties** toggle button can show or hide the panel that displays properties for selected objects. You can show and hide the following table using the **Show Table** toggle button. You can download the graphical timeline to SVG using the **Export to SVG** button.

3.17 Stable Diffusion Demo

Ryzen AI 1.8.0 provides preview demos of Stable Diffusion image-generation pipelines. The demos cover **Image-to-Image** (ControlNet: Canny, Pose, Tile, Depth; plus Segmind-Vega i2i without ControlNet) and **Text-to-Image** for SD1.5, SD-Turbo, SDXL-base, SDXL-Turbo, Segmind-Vega, DreamShaper XL Lightning, SSD-1B, Playground v2.5, FLUX.1-Schnell, FLUX.2-klein-4B, SD3.0, and SD3.5, with **inpainting** where noted for SD3.0.

Models are fetched from Hugging Face on first use (see *Installation Instructions* for Ryzen AI and driver setup). SD3.0 and SD3.5 models require Hugging Face authentication and might also require acceptance of the model license before they can be downloaded.

3.17.1 Installation Steps

1. Ensure the latest version of Ryzen AI and NPU drivers are installed. See *Installation Instructions*.
2. The GenAI-SD folder is located in the RyzenAI installation tree. Navigate to the folder and run the following command:

```
conda activate ryzen-ai-1.8.0
cd "$env:RYZEN_AI_INSTALLATION_PATH\GenAI-SD"
```

3. Models are downloaded from Hugging Face on first use and cached under `GenAI-SD\models` or `$env:RYZENAI_GENAI_SD_MODELS_ROOT`. See *supported-models* for the full list; primary entry points include:

    - SD1.5
    - SD1.5 ControlNet Canny
    - SD-Turbo
    - SDXL-Turbo
    - SDXL-base
    - Segmind-Vega
    - DreamShaper XL Lightning
    - SSD-1B
    - Playground v2.5 1024px
    - FLUX.1-Schnell
    - FLUX.2-klein-4B
    - SD3.0
    - SD3.5

3.17.2 Supported models

The following summarizes default application mode, typical resolution / dynamic-resolution (Dyn-Res) notes, and the recommended Hugging Face `model_id` on AMD NPU-tuned repos.

Notes	Model	App	Default resolution / DynRes	model_id (AMD Hub; stabilityai equivalent where applicable)
	SD1.5	t2i	512x512	amd/stable-diffusion-1.5-amdnpu
New	SD1.5	i2i-canny	512x512	amd/sd1.5-controlnet-canny-amdnpu
	SD-Turbo	t2i	512x512	amd/sd-turbo-amdnpu (or stabilityai/sd-turbo-amdnpu)
	SDXL-Turbo	t2i	512x512, 5x DynRes	amd/sdxl-turbo-amdnpu (or stabilityai/sdxl-turbo-amdnpu)
	SDXL-base	t2i	1024x1024, 20x DynRes	amd/sdxl-base-amdnpu (or stabilityai/sdxl-base-amdnpu)
	Segmind-Vega	t2i / i2i	1024x1024	amd/segmind-vega-amdnpu
New	DreamShaper XL Lightning	t2i	1024x1024, 20x DynRes	amd/dreamshaper-xl-lightning-amdnpu
New	SSD-1B	t2i	1024x1024, 20x DynRes	amd/SSD-1B-amdnpu
New	Playground v2.5	t2i	1024x1024, 20x DynRes	amd/playground-v2.5-1024px-aesthetic-amdnpu
New	FLUX.1-Schnell	t2i	1024x1024, 20x DynRes	amd/FLUX.1-schnell-amdnpu
New	FLUX.2-klein-4B	t2i	1024x1024	amd/FLUX.2-klein-4B-amdnpu
	SD3.0	t2i / ControlNet / Depth / Canny / Pose / Tile	512x512 (t2i default), 20x DynRes where applicable	stabilityai/stable-diffusion-3-medium-amdnpu
New	SD3.0	i2i-inpainting	1024x1024	stabilityai/stable-diffusion-3-medium-amdnpu
	SD3.5	t2i / ControlNet	512x512 (t2i default), 20x DynRes where applicable	stabilityai/stable-diffusion-3.5-medium-amdnpu
	SD3.5-ControlNet (Canny)	i2i-canny	512x512, 20x DynRes	Uses SD3.5 AMD weights plus SD3.0 Canny ControlNet assets (see below)

In the Supported models table, the **DynRes** column counts how many dynamic-resolution presets

a pipeline exposes (for example 5x or 20x). The following material spells out which width x height pairs are in scope.

Dynamic resolution (DynRes)

Original dynamic-resolution requirements

These are the classic fixed pairs (per model class) before the expanded preset grid.

- **DynRes 5x:** 512x512, 288x512, 512x288, 384x512, 512x384.
- **DynRes 20x:**

Table 8: New DynRes combinations (batch size 1; align to multiple of 32 as required)

Aspect (WxH)	Base 512x512	Base 640x640	Base 768x768	Base 1024x1024
1:1	512x512	640x640	768x768	1024x1024
4:3	512x384	640x480	768x576	1024x768
3:4	384x512	480x640	576x768	768x1024
16:9	512x288	640x384	768x448	1024x576
9:16	288x512	384x640	448x768	576x1024

Public model pages follow the Hugging Face model license (HF LIC) for each repo.

SD3.5 Canny ControlNet setup: copy the SD3.0 Canny ControlNet files into the SD3.5 model layout as required by your GenAI-SD tree (per release notes), then run the Canny example with `--model_id amd/stable-diffusion-3.5-medium-amdnpu`.

3.17.3 Running the Demos

Activate the conda environment (see also **`installation-steps`_**):

```
conda activate ryzen-ai-1.8.0
```

Optionally, set the NPU to high performance mode to maximize performance:

```
xrt-smi configure --pmode performance
```

Refer to the documentation on *xrt-smi configure* for additional options.

From the `GenAI-SD\test` directory unless noted otherwise. All examples use the unified entry point `run.py` and pass `--model_id` with the Hugging Face model identifier. For more details, refer to the `README.md` file in the `GenAI-SD` directory.

Image-to-Image with ControlNet

The image-to-image demo generates images from a **prompt** plus a **control image** (ControlNet types such as Canny, pose, tile, or depth for SD3.x, selected with `-C`). SD3.x often defaults to 512x512; override resolution with `-W` and `-H` as in the examples below. SD3.x DynRes presets are summarized in *supported-models* and Dynamic resolution (DynRes).

To run a minimal Canny example:

```
python run.py -C canny --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu
```

The demo can use `.\ref\canny.jpg` as the control image unless you override `--control_image_path`. Outputs are written to `generated_images` by default, unless `--output_path` is specified. You can redirect the output directory to a location where your user account has write permissions, for example: `--output_path C:\Users\<username>\Documents\generated_images`.

SD1.5 ControlNet Canny (i2i-canny)

```
python .\run.py --model_id amd/sd1.5-controlnet-canny-amdnpu
```

Segmind-Vega (i2i, no ControlNet path)

```
python .\run.py --model_id amd/segmind-vega-amdnpu --control_image_path .\
↪assets\controlimg_input_1024x1024.png --strength 0.95
```

SD3.0 ControlNet (canny / pose / tile / depth)

```
python .\run.py -C canny --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu --prompt "Anime style illustration of a girl wearing a suit. A␣
↪moon in sky. In the background we see heavy rain approaching. text
↪'InstantX' on image" -H 1024 -W 1024 --control_image_path .\ref\canny.
↪jpg -n 50
python .\run.py -C pose --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu --prompt "Anime style illustration of a girl wearing a suit. A␣
↪moon in sky. In the background we see heavy rain approaching. text
↪'InstantX' on image" -H 1024 -W 1024 --control_image_path .\ref\pose.
↪jpg -n 50
python .\run.py -C tile --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu --prompt "Anime style illustration of a girl wearing a suit. A␣
↪moon in sky. In the background we see heavy rain approaching. text
↪'InstantX' on image" -H 1024 -W 1024 --control_image_path .\ref\tile.
↪jpg -n 50
python .\run.py -C depth --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu -H 1024 -W 1024 --control_image_path .\assets\depth.jpeg -n 50
```

SD3.5 ControlNet Canny (i2i-canny)

After copying SD3.0 Canny ControlNet into the SD3.5 layout as required:

```
python .\run.py -C canny --model_id stabilityai/stable-diffusion-3.5-
↪medium-amdnpu --prompt "Anime style illustration of a girl wearing a␣
↪suit. A moon in sky. In the background we see heavy rain approaching.␣
↪text 'InstantX' on image" -H 1024 -W 1024 --control_image_path .\ref\
↪canny.jpg -n 50
```

For SD3 and SD3.5 models, pass `-O1` (or `--optimize_o1`) to enable a performance optimization preset that skips selected DiT denoising steps to improve inference performance.

Text-to-Image

The text-to-image demo generates images from **text prompts only** (no control image). It covers SD1.5 (512-class), SD-Turbo and SDXL-Turbo (512-class), SDXL-base, Segmind-Vega, DreamShaper XL Lightning, SSD-1B, Playground v2.5, FLUX.1-Schnell, FLUX.2-klein-4B, and SD3.0 / SD3.5 with `-C None`. Use `-H`, `-W`, and `-n` when the pipeline supports them.

```
python .\run.py --model_id amd/stable-diffusion-1.5-amdnpu
python .\run.py --model_id stabilityai/sd-turbo-amdnpu
python .\run.py --model_id stabilityai/sdxl-turbo-amdnpu
python .\run.py --model_id stabilityai/sdxl-base-amdnpu
python .\run.py --model_id amd/segmind-vega-amdnpu
python .\run.py --model_id amd/dreamshaper-xl-lightning-amdnpu
python .\run.py --model_id amd/SSD-1B-amdnpu
python .\run.py --model_id amd/playground-v2.5-1024px-aesthetic-amdnpu
python .\run.py --model_id amd/FLUX.1-schnell-amdnpu
python .\run.py --model_id amd/FLUX.2-klein-4B-amdnpu
python .\run.py -C None --model_id stabilityai/stable-diffusion-3-medium-
↪amdnpu -H 1024 -W 1024 -n 50
python .\run.py -C None --model_id stabilityai/stable-diffusion-3.5-
↪medium-amdnpu -H 1024 -W 1024 -n 50
```

Custom prompts can be supplied with `--prompt`. For example:

```
python .\run.py --model_id stabilityai/stable-diffusion-1.5-amdnpu --
↪prompt "Photo of an ultra realistic sailing ship, dramatic light, pale␣
↪sunrise, cinematic lighting, battered, low angle, trending on␣
↪artstation, 4k, hyper realistic, focused, extreme details"
```

Inpainting

SD3.0 (``**-C Inpainting**``) uses a base image and `--control_mask_path` (URLs or local paths). This extends the image-conditioned flows above with an explicit mask channel.

```
python .\run.py --model_id stabilityai/stable-diffusion-3-medium-amdnpu -
↪C Inpainting --prompt "A cat is sitting next to a puppy" --n_prompt
↪"deformed, distorted, disfigured, poorly drawn, bad anatomy, wrong␣
↪anatomy, extra limb, missing limb, floating limbs, mutated hands and␣
↪fingers, disconnected limbs, mutation, mutated, ugly, disgusting,␣
↪blurry, amputation, NSFW" -n 28 --controlnet_conditioning_scale 0.95 --
↪control_image_path "https://huggingface.co/alimama-creative/SD3-
↪Controlnet-Inpainting/resolve/main/images/dog.png" --control_mask_path
↪"https://huggingface.co/alimama-creative/SD3-Controlnet-Inpainting/
↪resolve/main/images/dog_mask.png" --seed 42
```

3.17.4 Running with AMD Stable Diffusion Sandbox

AMD SD Sandbox is a framework for running Stable Diffusion (SD) models accelerated by AMD Ryzen AI hardware. It provides an easy-to-use interface for evaluating, comparing, and deploying multiple SD pipelines.

Please go to the AMD SD Sandbox GitHub repository for more information.

3.18 Ryzen AI CVML library

The Ryzen AI Libraries build on top of the Ryzen AI drivers and execution infrastructure to provide powerful AI capabilities to C++ applications without having to worry about training specific AI models and integrating them to the Ryzen AI framework.

Each Ryzen AI CVML library feature offers a simple C++ application programming interface (API) that can be easily incorporated into existing applications.

Table of Contents

- *Package Contents*
- *Executing Ryzen AI CVML library enabled applications*
  - *Prerequisites and dependencies*
  - *Windows Setup*
  - *Ubuntu Setup*
- *Programming guide for C++ Applications*

  - *Include Ryzen AI CVML library headers*
  - *Create Ryzen AI CVML library context*
  - *Create Ryzen AI CVML library feature object*
  - *Encapsulate image buffers*
  - *Execute the feature*
- *Building applications with Ryzen AI CVML Libraries*
  - *Building Ryzen AI CVML library applications with CMake*
  - *Building Ryzen AI CVML library sample applications*
- *Running Ryzen AI CVML library sample applications*
  - *Locating Ryzen AI CVML library runtime files*
  - *Select an input source/image/video*
  - *Execute the sample application*

3.18.1 Package Contents

The Ryzen AI CVML library consists of the following files and folders:

- **cmake/** — Packaging info for CMake's find_package function
- **include/** — C++ header files
- **windows/** — Binary files for Windows, including both compile time .LIB files and runtime .DLL/.GRAPHLIB/.AMODEL files
- **linux/** — Binary files for Linux, including compile and runtime .SO files
- **samples/** — Individual sample applications
- **LICENSE.txt** — License file

3.18.2 Executing Ryzen AI CVML library enabled applications

The Ryzen AI CVML library selects the appropriate hardware (e.g., GPU or NPU) and framework for performing inference operations by default. An API is also available to set the preferred inference backend for those applications that wish to do so.

In order to execute applications that utilize the Ryzen AI CVML library, the appropriate drivers must first be installed on the target system, and the Ryzen AI CVML library files must be included with the application itself.

> **Note**
>
> Ryzen AI CVML library features that utilize the ONNX backend for NPU operations might experience a longer startup latency the first time they are executed on a device. This increased startup latency does not occur for subsequent runs of the feature.

> **Note**
>
> If the NPU driver is not installed on the target system, the Ryzen AI CVML library will automatically fall back to the GPU backend for inference operations.

Prerequisites and dependencies

The AMD Adrenalin and Ryzen AI drivers should be installed before attempting to execute Ryzen AI CVML library applications.

Download Ryzen AI CVML Library package

Create an AMD account at account.amd.com if you don't have one, then sign in to download the Ryzen AI CVML Library from the AMD Account Portal:

```
https://account.amd.com/en/forms/downloads/xef.html?filename=72293_Ryzen_
↪AI_Library_26.07.15.zip
```

After downloading, extract the package to a local directory (e.g., `C:\RyzenAI-Library` on Windows or `~/RyzenAI-Library` on Linux) and set the `AMD_CVML_SDK_ROOT` environment variable to the extracted location. s .. _windows_setup:

Windows Setup

The following installations are for Windows OS. For Linux OS, follow *Ubuntu Setup* instructions.

AMD Adrenalin driver

Install either the following Adrenalin driver or a newer one: https://www.amd.com/en/support/download/drivers.html

AMD Ryzen AI driver

Install the latest Ryzen AI NPU driver from https://ryzenai.docs.amd.com/en/latest/inst.html

Version 32.0.203.280 or newer is required; both listed versions (32.0.203.280 and 32.0.203.314) are compatible.

OpenCV

Download the OpenCV 4.11 Windows installer from the GitHub releases page: https://github.com/opencv/opencv/releases/tag/4.11.0

Download `opencv-4.11.0-windows.exe`, run it, and extract to a local folder (e.g. `C:\opencv`). The `build` subfolder (e.g. `C:\opencv\build`) contains `OpenCVConfig.cmake` and is the path to use for `OPENCV_INSTALL_ROOT`.

Ubuntu Setup

The following installations are for Ubuntu. For Windows OS, follow *Windows Setup* instructions.

Ensure that the following software tools/packages are installed on the development system:

1. OS: Ubuntu 22.04 or Ubuntu 24.04 (linux kernel >= 6.11.0-21-generic)
2. Install latest Ryzen AI NPU driver following the "Install NPU Drivers" section
3. Vulkan SDK
4. OpenCV 4.11.0 — build from source following the instructions in `README-linux.md`

Installing VulkanSDK and 22.04/24.04 specific installs

```
UBUNTU_CODENAME=$(. /etc/os-release; echo "$UBUNTU_CODENAME")
wget -qO- https://packages.lunarg.com/lunarg-signing-key-pub.asc | sudo␣
↪tee /etc/apt/trusted.gpg.d/lunarg.asc
sudo wget -qO /etc/apt/sources.list.d/lunarg-vulkan-1.3.296-$UBUNTU_
↪CODENAME.list https://packages.lunarg.com/vulkan/1.3.296/lunarg-vulkan-
↪1.3.296-$UBUNTU_CODENAME.list
sudo apt update
sudo apt install vulkan-sdk
```

Additional installation for Ubuntu 22.04: Update MESA Vulkan Drivers

```
sudo apt update && sudo apt upgrade
sudo add-apt-repository ppa:kisak/kisak-mesa -y
sudo apt update
sudo apt upgrade
```

Additional installation for Ubuntu 24.04

```
sudo apt install libavcodec-dev libavformat-dev libswscale-dev libnsl2␣
↪gstreamer1.0-plugins-good gstreamer1.0-plugins-bad gstreamer1.0-plugins-
↪ugly -y
```

(continues on next page)

(continued from previous page)

```
DEP_PKG_LIST="https://launchpad.net/ubuntu/+archive/primary/+files/
↪libmpdec3_2.5.1-2build2_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libpython3.10-
↪minimal_3.10.4-3_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libpython3.10-
↪stdlib_3.10.4-3_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libpython3.10_3.
↪10.4-3_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libprotobuf23_3.
↪12.4-1ubuntu7_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libgoogle-
↪glog0v5_0.5.0+really0.4.0-2_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libtiff5_4.3.0-6_
↪amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libilmbase25_2.5.
↪7-2_amd64.deb \
    https://launchpad.net/ubuntu/+archive/primary/+files/libopenexr25_2.5.
↪7-1_amd64.deb"

for pkg in $DEP_PKG_LIST
do
    echo $pkg
    wget $pkg
    sudo dpkg -i *.deb
    rm *.deb
done
```

3.18.3 Programming guide for C++ Applications

Incorporating the Ryzen AI's optimized features into C++ applications can be done in a few simple steps, as explained in the following sections.

Include Ryzen AI CVML library headers

The required definitions for compiling each Ryzen AI feature are included in a corresponding header file under the **include/** folder:

```
cvml-feature-name.h
```

where `feature-name` is the name of the desired Ryzen AI feature.

For example, the definitions for the Ryzen AI Depth Estimation feature are available after adding a line similar to the following example:

```
#include <cvml-depth-estimation.h>
```

Details about each feature's programming interface and expected usage are provided within their individual include headers.

Create Ryzen AI CVML library context

Each Ryzen AI CVML library feature is created against a *CVML context* (see `amd::cvml::Context`). The context provides access to common functions for logging, etc. A pointer to a new context may be obtained by calling the `amd::cvml::CreateContext()` function:

```
auto ryzenai_context = amd::cvml::CreateContext();
```

When no longer needed, the context may be released using its `Release()` member function:

```
ryzenai_context->Release();
```

Create Ryzen AI CVML library feature object

The application programming interface for each feature is provided via a *Ryzen AI CVML library C++ feature object* that may be instantiated once a Ryzen AI CVML library context has been created.

The following example instantiates a feature object for the depth estimation library:

```
amd::cvml::DepthEstimation ryzenai_depth_estimation(ryzenai_context);
```

Encapsulate image buffers

The Ryzen AI CVML library defines its own *Image* class (see `amd::cvml::Image`) for representing images and video frame buffers. Each *Image* object is assigned a specific format and data type on creation. For example, an *Image* to encapsulate an incoming RGB888 frame buffer can be created with the following code:

```
amd::cvml::Image ryzenai_image(amd::cvml::Image::Format::kRGB,
                               amd::cvml::Image::DataType::kUint8, width,
                               height, data_pointer);
```

Execute the feature

To execute a Ryzen AI feature on a provided input, call the appropriate *execution* member function of the Ryzen AI CVML library feature object.

For example, the following code executes a single instance of the depth estimation library, using the *ryzenai_image* from the previous section:

```
// encapsulate output buffer
amd::cvml::Image ryzenai_output(amd::cvml::Image::Format::kGrayScale,
                                amd::cvml::Image::DataType::kFloat32,
                                output_width, output_height, output_
↪pointer);

// execute the feature
ryzenai_depth_estimation.GenerateDepthMap(ryzenai_image, &ryzenai_output);
```

3.18.4 Building applications with Ryzen AI CVML Libraries

When building applications against the Ryzen AI CVML library, ensure that the library's `include/` folder is part of the compiler's include paths, and that the library's `windows/` or `linux/` folder has been added to the linker's library paths.

Depending on the application's build environment, it may also be necessary to explicitly list which of the Ryzen AI CVML library's .LIB files (when building for Windows applications) need to be linked.

Building Ryzen AI CVML library applications with CMake

If CMake is used for the application's build environment, the necessary include folder and link libraries can be added with the following lines in the application's `CMakeLists.txt` file:

```
# find Ryzen AI CVML library and set include folders
find_package(RyzenAILibrary REQUIRED PATHS ${AMD_CVML_SDK_ROOT})

# add Ryzen AI CVML library linker libraries
target_link_libraries(${PROJECT_NAME} ${RyzenAILibrary_LIBS})
```

where `AMD_CVML_SDK_ROOT` defines the location of the Ryzen AI CVML library files and `PROJECT_NAME` defines the name of the application build target.

Building Ryzen AI CVML library sample applications

In addition to general Ryzen AI CVML library prerequisite and dependencies listed under *Prerequisites and dependencies*, the included sample applications also make use of OpenCV for reading input images/videos/camera and displaying final output windows. A copy of OpenCV will need to be downloaded to the development system before the samples can be rebuilt and/or executed. Note that CVML samples are built and tested with OpenCV 4.11 on both Windows and Linux.

Ensure the following prerequisites have been set up to build Ryzen AI CVML library sample applications:

- CMake has been installed and is available in the system/user path

- On Windows, Visual Studio's "Desktop development with C++" build tools, or a comparable C++ toolchain, has been installed
- The location of OpenCV libraries has been assigned to the `OPENCV_INSTALL_ROOT` environment variable
- The relative locations of the `include`, `windows`, `linux`, and `samples` folders are unchanged

The following are CMake commands for building samples.

On Windows,

```
rem Point to the build subfolder inside your OpenCV installation
rem (e.g. if you extracted OpenCV to C:\opencv, use C:\opencv\build)
rem CMake's find_package needs this folder to locate OpenCVConfig.cmake
set OPENCV_INSTALL_ROOT=C:\opencv\build
cd samples/
mkdir build
cmake -S %CD% -B %CD%\build -DOPENCV_INSTALL_ROOT=%OPENCV_INSTALL_ROOT% -
↪DCMAKE_PREFIX_PATH=%OPENCV_INSTALL_ROOT%
cmake --build %CD%\build --config Release
```

On Linux,

```
export OPENCV_INSTALL_ROOT=<Path to OpenCV Libraries>
cd samples/
mkdir build
cmake -S $PWD -B $PWD/build -DOPENCV_INSTALL_ROOT="$OPENCV_INSTALL_ROOT" -
↪DCMAKE_PREFIX_PATH="$OPENCV_INSTALL_ROOT"
cmake --build $PWD/build --config Release
```

3.18.5 Running Ryzen AI CVML library sample applications

This section describes how to execute Ryzen AI CVML library sample applications.

Locating Ryzen AI CVML library runtime files

When executing applications built against the Ryzen AI CVML library, ensure that the runtime files are accessible as described below for each platform. For both Windows and Linux, add OpenCV runtime libs to PATH.

For Windows, either one of the following conditions are met:

1. The Ryzen AI CVML library runtime dll and graphlib files are in the same folder as the application executable.
2. The Ryzen AI CVML library's **windows/** folder has been added to the PATH environment variable.

```
set PATH=<location of Ryzen AI CVML library package>\windows;%PATH%
set PATH=%OPENCV_INSTALL_ROOT%\x64\vc16\bin;%PATH%
```

For Linux, all the following conditions are met:

1. Add location of **linux/** to LD_LIBRARY_PATH
2. Add location of NPU driver libs (`/opt/xilinx/xrt/lib`) to LD_LIBRARY_PATH

```
export LD_LIBRARY_PATH=<location of Ryzen AI CVML library package>/linux:
↪$LD_LIBRARY_PATH
export LD_LIBRARY_PATH=/opt/xilinx/xrt/lib:$LD_LIBRARY_PATH
export LD_LIBRARY_PATH=$OPENCV_INSTALL_ROOT/lib:$LD_LIBRARY_PATH
```

Select an input source/image/video

Ryzen AI CVML library samples can accept a variety of image and video input formats, or even open the default camera on the system if "0" is specified as an input.

In this example, a publicly available video file is used for the application's input.

```
curl -o dancing.mp4 https://videos.pexels.com/video-files/4540332/4540332-
↪hd_1920_1080_25fps.mp4
```

Execute the sample application

Finally, the previously built sample application can be executed with the selected input source.

On Windows,

```
build\cvml-sample-depth-estimation\Release\cvml-sample-depth-estimation.
↪exe -i dancing.mp4
```

On Linux,

```
./build/cvml-sample-depth-estimation/cvml-sample-depth-estimation -i␣
↪dancing.mp4
```

3.19 Supported Operators

The following table lists the ONNX operators supported in Ryzen AI 1.8.0. The operators are categorized by their support for different quantization types: BF16, A16W8, A8W8, and XINT8. A "Y" indicates that Ryzen AI provides broad coverage for that operator and for that specific quantization type for CNN and NLP models. Some specific configurations of that operator may however not be fully supported.

Ops	BF16	A16W8	A8W8	XINT8
Abs	Y			Y
Add	Y	Y	Y	Y
And	Y			Y
ArgMax	Y			Y
ArgMin	Y			Y
AveragePool	Y	Y	Y	Y
BatchNormalization	Y	Y	Y	Y
BitShift	Y			Y
BitwiseAnd	Y			Y
BitwiseNot	Y			Y
BitwiseOr	Y			Y
BitwiseXor	Y			Y
Cast	Y			Y
Ceil	Y			Y
Celu				Y
Clip		Y	Y	Y
Concat	Y	Y	Y	Y
Constant	Y	Y	Y	Y
ConstantOfShape	Y	Y	Y	Y
Conv	Y	Y	Y	Y
ConvTranspose	Y	Y	Y	
Cos		Y	Y	
CumSum	Y			Y
DepthToSpace	Y	Y	Y	Y
DequantizeLinear		Y	Y	Y
Div	Y	Y	Y	Y
Einsum	Y			
Elu		Y		Y
Equal	Y			Y
Erf	Y			Y
Exp	Y	Y		
Expand	Y	Y	Y	Y
Flatten	Y	Y	Y	Y
Floor	Y			Y
Gather	Y			
GatherElements	Y	Y		
Gelu		Y	Y	Y
Gemm	Y	Y	Y	Y
GlobalAveragePool	Y	Y	Y	Y
GlobalMaxPool		Y	Y	Y
Greater	Y			Y

continues on next page

Table 9 – continued from previous page

Ops	BF16	A16W8	A8W8	XINT8
GreaterOrEqual				Y
GridSample	Y			
GroupConv	Y			
GroupNormalization		Y	Y	Y
HardSigmoid			Y	
HardSwish			Y	Y
Identity	Y	Y	Y	Y
InstanceNormalization	Y	Y	Y	Y
LSTM	Y			
LayerNormalization				Y
LeakyRelu		Y	Y	Y
Less	Y			Y
LessOrEqual				Y
Log	Y			Y
MatMul	Y	Y	Y	Y
Max	Y	Y	Y	Y
MaxPool	Y	Y	Y	Y
Min	Y	Y	Y	Y
Mish				Y
Mod	Y			
Mul	Y	Y	Y	Y
Neg	Y			Y
Not	Y			
Or	Y			Y
Pad	Y	Y	Y	Y
Pow	Y			Y
QLinearConv		Y	Y	
QLinearMatMul			Y	
QuantizeLinear		Y	Y	
Range				Y
Reciprocal	Y	Y	Y	Y
ReduceMax	Y	Y	Y	Y
ReduceMean	Y		Y	Y
ReduceMin	Y		Y	Y
ReduceSum	Y		Y	Y
Relu		Y	Y	Y
Reshape	Y	Y	Y	Y
Resize	Y		Y	Y
Round	Y			Y
STFT				Y
ScatterND	Y			

continues on next page

Table 9 – continued from previous page

Ops	BF16	A16W8	A8W8	XINT8
Shape	Y	Y	Y	Y
Shrink				Y
Sigmoid	Y	Y	Y	
Sign	Y			Y
Sin	Y	Y	Y	
Size	Y			
Slice	Y	Y	Y	Y
Softmax		Y	Y	Y
Softsign				Y
Split	Y	Y		
Sqrt	Y	Y	Y	Y
Squeeze	Y	Y	Y	Y
Sub	Y	Y	Y	Y
Tanh	Y	Y		Y
ThresholdedRelu				Y
Tile	Y			Y
TopK		Y		Y
Transpose	Y	Y	Y	Y
Unsqueeze	Y	Y	Y	Y
Upsample		Y	Y	Y
Upsample (deprecated)	Y			
Where	Y	Y		Y
Xor	Y			

3.19.1 LLM Operator support

The Ryzen AI LLM execution flow supports ONNX Runtime GenAI-based models with the following operators:

- SimplifiedLayerNormalization
- SkipSimplifiedLayerNormalization
- MatMulNBits (W4ABF16/W4ABFP16)
- Add
- RotaryEmbedding
- GroupQueryAttention
- Sigmoid
- Mul

These operators are optimized for execution on the Ryzen AI NPU via the ONNX Runtime GenAI framework.

3.20 Licensing Information

Ryzen AI is released by Advanced Micro Devices, Inc. (AMD) and is subject to the licensing terms listed below. Some components may include third-party software that is subject to additional licenses. Review the following links for more information:

- [AMD End User License Agreement](https://account.amd.com/content/dam/account/en/licenses/download/ryzen-ai-end-user-license-agreement-public.pdf)
- [Windows - Third Party End User License Agreement](https://account.amd.com/content/dam/account/en/licenses/download/ryzen-ai-1-8-0-ga-tpn-license.pdf)
- [Linux - Third Party End User License Agreement](https://account.amd.com/content/dam/account/en/licenses/download/ryzen-ai-1-8-0-linux-tpn-license.pdf)