



A GSTREAMER- BASED PERFORMANCE EVALUATION TOOL FOR END-TO-END VISUAL AI PIPELINES

JAN GORECKI, INTEL

OCT 11, 2026 - GSTREAMER CONFERENCE

AGENDA



01 Problem Statement – how to evaluate/size performance of my AI System?

02 What are the options?


03 How to run a performance evaluation for Gstreamer based AI Pipeline - Demo



01 | PROBLEM STATEMENT




Problem Statement



- What is performance of my AI System?
- Is my hardware / system sufficient for the use case I am running?
- Am I overdimensioning / underdimensioning my AI System?
- How to evaluate my system quickly?
- What advise to give to my customer?



What are variables?



There are
many
variables and
some more...

Why TOPS Alone Is Insufficient for Platform Evaluation






Image Complexity

Codec • Resolution • Bit rate • Number of objects




Pipeline Complexity

Motion detector • # AI models • Media workload •
Other post-processing



Model Complexity

Model architecture • Quantization



System Resource

CPU / GPU / NPU • Memory Speed • Power budget



Workload Optimization

Workload placement • Batch size • Inference interval





02 | WHAT ARE THE OPTIONS?



What are resolution options?

- Access pre-benchmarked results that are published
 - E.g. [OpenVINO benchmarks](#)
- Use Gstreamer based (DLStreamer) evaluation tool
[Visual Pipeline and Platform Evaluation Tool](#)



Targeted Personas



Hardware Purchasing Decision Maker


- Decide which processor and accelerator needed for job
- Want to benchmark platform with a sample workload.
- Benchmarks different configurations to answer what if questions.

Platform Marketing Engineer / Consulting Engineer

- Understand platform potential
- Want a tool to benchmark the platform.
- Want to demonstrate platform capabilities to customers



Easy End-to-End Visual AI Benchmarking with ViPPET




About ViPPET

- Evaluates end-to-end platform performance and capacity for real-world visual AI workloads
- Enables rapid “what-if” analysis across models, pipelines and system configurations
- Provide system-level insights beyond raw TOPs by measuring full video + AI workloads

Recommended Usage

- Assess platform performance and scaling behavior using proxy pipelines and widely adopted AI models
- Simple, intuitive UI with real-time telemetry charts for clear, customer-facing performance demonstrations



Predefined Pipelines to get started



- Smart Parking
- Smart NVR
- Simple NVR
- Segmentation
- Motion Detection
- License Plate Recognition
- Human Pose Detection
- Age & Gender Detection
- Goods Detection
- Goods Detection & Classification

.. and more to come

PERFORMANCE RUN CONFIGURATION



Throughput Test

- Measures average FPS
- Multiple pipelines simultaneously

Density Test


- Determines platform ceiling for a given throughput target



ViPPET - Visual Pipeline and Platform Evaluation Tool

🏠 > Manufacturing AI Suite > Visual Pipeline and Platform Evaluation Tool

- Opensourced, Apache 2.0 license
- Quarterly releases
- [Documentation Page](#) 
- [Github page](#) 





Visual Pipeline and Platform Evaluation Tool

What is ViPPET

The Visual Pipeline and Platform Evaluation Tool simplifies hardware selection for AI workloads by enabling configuration of workload parameters, performance benchmarking, and analysis of key metrics such as throughput, CPU usage, and GPU usage. With its intuitive interface, the tool provides actionable insights that support optimized hardware selection and performance tuning.

ViPPET helps you compare how AI pipelines behave across hardware targets, model variants, and input sources so you can make better deployment decisions earlier in the development cycle.






ViPPET flow and functionalities


1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration
- Camera input



2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)



3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing



4 Enhanced User Experience & Visualization

- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode





03 | VIPPET DEMO



Takeaways and questions

- ViPPET is released with quarterly cadence
- End-to-end workload performance matters
- ViPPET helps size the hardware by measuring pipeline performance and stream capacity
- Thank you - contact: jan.gorecki@intel.com

ViPPET – new features

Visual Pipeline and Platform Evaluation Tool

<https://docs.openedgeplatform.intel.com/2026.2/edge-ai-libraries/visual-pipeline-and-platform-evaluation-tool/index.html>

ViPPET feature structure

1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration
- Camera input

2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)

3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing

4 Enhanced User Experience & Visualization

- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode

ViPPET feature structure

1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration

2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)

3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing


4 Enhanced User Experience & Visualization

- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode


Pipeline Templates

Predefined and Highly Optimized Detect and Detect + Classify Pipeline Templates as Starting Points for Creating New Pipelines

- Challenge
 - In ViPPET 2025.2, the only option to create new pipelines was from uploading GStreamer descriptions (gst-launch cmd).
- Solution
 - This release introduces two predefined pipeline templates as starting points for creating new pipelines
 - „Detect” template
 - „Detect & Classify” template
 - These templates are based on M9 pipelines and their BKM.
- Benefits
 - Simple way to create new pipelines
 - New pipelines will be optimized from the get-go



„Detect” template




„Detect & Classify” template

Camera Support

Support for USB and Network Cameras

- Challenge
 - ViPPET 2025.2 offered videos as the only option for input source for pipelines.
- Solution
 - This release introduces the option to use cameras as an input source for pipelines.
 - Camera discovery service automatically detects both USB and network cameras (using ONVIF), which can be used as the input sources in the pipeline editor.
- Benefits
 - It is now possible to run ViPPET pipelines using connected USB and network cameras.
 - Faster setup and deployment; stronger demos and presentations.



Video Upload Support

Users can now upload their own video files and use them as input for pipelines

- Challenge
 - ViPPET 2026.0 offered only a few predefined videos as input for pipelines
- Solution
 - Users can now upload their own video files to ViPPET via the Videos panel, which provides a file upload control with progress, success confirmation, and error handling.
 - Uploaded videos are added to the list of available inputs and can be selected directly in the Pipeline Builder. The maximum supported file size is 2 GB.
- Benefits
 - Users can now upload their own video files and use them as input for pipelines

Videos

Ready-to-use video clips available in the platform

Upload Files



Drag & drop files here
or click to browse your computer



File name	Resolution	Number of frames	Codec	Duration	
bears.mp4	3840x2160	1080	h264	0m 18s	
metro_smart_parking.mp4	640x480	3000	h264	1m 40s	
obj_classification.mp4	3840x2160	1715	h264	0m 57s	
people.mp4	1920x1080	302	h264	0m 10s	

Image-Set Upload Support

Users can now upload image files to ViPPET and use them as input for pipelines

- Challenge
 - ViPPET 2026.0 offered only a few predefined videos as input for pipelines
- Solution
 - Users can now upload image files to ViPPET as input for pipelines. Images can be uploaded as an archive (.zip, .tar, .tar.gz, .tgz) with specific structure (sequentially named files).
 - Images are played back using GStreamer *multifilesrc*.
- Benefits
 - Users can now upload image files to ViPPET and use them as input for pipelines



Custom Model Upload Support

Users can now upload OpenVINO models, including models trained using Intel Geti platform

- Challenge
 - ViPPET 2026.0 offered only a few predefined models to use for pipelines
- Solution
 - Users can now upload OpenVINO models, including models trained using the Intel Geti platform.
 - Uploaded models are added to the list of available models in ViPPET, making them easier to manage and reuse across pipelines.
- Benefits
 - This feature simplifies bringing your own models into the application for testing and deployment.

The screenshot displays the 'Models' section of the ViPPET interface. At the top, it says 'Ready-to-use models available in the platform'. Below this is an 'Upload Files' section with a dashed box and an upload icon, with the text 'Drag & drop files here or click to browse your computer'. To the right of the upload box, it says '1 model selected' and 'Install (1)'. Below the upload box is a table of models:

☒	Name	Category	Source	Status	Precisions	Used by	Actions
☐	PaddleOCR	classification	ultralytics	Installed	FP32	License Plate Recognition	
☐	EfficientNet B0	classification	pipeline-zoo-models	Installed	INT8	Goods Detection & Classification	
☐	MobileNet V2 PyTorch	classification	omz	Installed	FP16	Smart NVR	
☐	Age Gender Recognition Retail 0013	classification					
☐	Color Classification	classification					
☐	YOLO v8 License Plate Detector	detection					
☒	YOLO v8n 640x640	detection					

Below the table is another 'Models' section with the same header. Below this is another 'Upload Files' section with a dashed box and an upload icon, with the text 'Drag & drop files here or click to browse your computer'. Below the upload box is a form with the following fields:

- Model name:
- Category:
- Selected files (1):

OpenVINO_model.zip 7.94 MB ×
- Buttons:

ViPPET feature structure

1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration

2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)

3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing

4 Enhanced User Experience & Visualization

- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode

New Predefined Pipelines (1)

Predefined Pipelines for Retail, Manufacturing, and Metro Scenarios

- Challenge
 - ViPPET 2025.2 offered limited set of predefined pipelines (i.e. NVR & ALPR)
- Solution
 - New **Retail** pipelines: Goods Detection, Goods Detection & Classification, Age & Gender Recognition
 - New **Manufacturing** pipelines: Defect Detection
 - New **Metro** pipelines: Smart Parking
- Benefits
 - End-to-end solutions for more out-of-the-box vertical business use cases

The image displays a grid of eight predefined pipeline cards, each representing a different AI solution. Each card includes a header image, a title, a category and hardware selection bar, and a brief description.

- Smart Parking**: Category: Smart Cities. Hardware: CPU, GPU, NPU, GPU_NPU. Description: Smart Parking is a cloud-native, microservices-based video analytics pipeline that uses pre-trained deep learning models to detect parking-space occupancy.
- Smart NVR**: Category: Smart Cities. Hardware: CPU, GPU, NPU, GPU_NPU. Description: This pipeline processes video streams with AI-based analytics including object detection, tracking, and classification. Outputs include video recording, metadata, and processed video frames.
- Simple NVR**: Category: Smart Cities. Hardware: CPU, GPU, NPU, GPU_NPU. Description: This pipeline performs basic video decoding, recording, and format conversion for lightweight media handling.
- License Plate Recognition**: Category: Smart Cities. Hardware: CPU, GPU, NPU, GPU_NPU. Description: The Simple Video Structurization (D-T-C) pipeline is a versatile, use case-agnostic solution that supports license plate recognition, vehicle detection with attribute classification, and other object detection and classification tasks, adaptable based on...
- Goods Detection**: Category: Retail. Hardware: CPU, GPU, NPU, GPU_NPU. Description: The Retail Pipeline using YOLO 11n for object detection is designed for retail analytics applications. It detects various retail-related objects, making it suitable for inventory management, customer behavior analysis, and other retail use cases.
- Goods Detection & Classification**: Category: Retail. Hardware: CPU, GPU, NPU, GPU_NPU. Description: The Retail Pipeline using YOLO 11n for object detection and EfficientNet B0 for classification is designed for retail analytics applications. It detects various retail-related objects and classifies them into specific categories, making it suitable for inventory...
- Defect Detection**: Category: Manufacturing. Hardware: CPU, GPU, NPU, GPU_NPU. Description: AI-powered pipeline for automatic pallet defect detection using machine vision, enabling efficient quality control in manufacturing.
- Age & Gender Recognition**: Category: Retail. Hardware: CPU, GPU, NPU, GPU_NPU. Description: The Retail Pipeline using face-detection-retail-0004 for object detection and age-gender-recognition-retail-0013 for age and gender prediction is designed for retail analytics applications. It detects faces, predicts age groups and gender, making...

New predefined pipelines (1)

Four predefined pipelines showcasing Video Summarization, Motion Detection, Instance Segmentation, and Pose Estimation

- Challenge
 - ViPPET 2026.0 offered limited set of predefined pipelines
- Solution
 - Video Summarization: Video summarization pipeline using *gvagenai* with a vision-language model to generate concise scene-level summaries from sampled frames.
 - Motion Detection: Motion detection pipeline that uses *gvamotiondetect* to identify regions of motion, then runs YOLOv8n object detection restricted to those motion ROIs via *gvadetect*.
 - Instance Segmentation: Pipeline based on YOLO26n-seg, which involves identifying individual objects in an image and segmenting them from the rest of the image.
 - Pose Estimation: Pipeline based on YOLO26n-pose, which involves identifying the location of specific human body keypoints (eyes, ears, ...).
- Benefits
 - End-to-end solutions for more out of the box vertical business use cases


Pipeline	Description	Variants
	Video Summarization: Video summarization pipeline using <i>gvagenai</i> with a vision-language model to generate concise scene-level summaries from sampled frames.	Available in CPU, GPU and NPU variants.
	Motion Detection: Motion detection pipeline that uses <i>gvamotiondetect</i> to identify regions of motion, then runs YOLOv8n object detection restricted to those motion ROIs via <i>gvadetect</i> .	Available in CPU, GPU and NPU variants.
	Instance Segmentation: Pipeline based on YOLO26n-seg, which involves identifying individual objects in an image and segmenting them from the rest of the image.	Available in CPU and GPU variants.
	Pose Estimation: Pipeline based on YOLO26n-pose, which involves identifying the location of specific human body keypoints (eyes, ears, ...).	Available in CPU, GPU and NPU variants.

New predefined pipelines (2)

New pipelines for showcasing DeepSORT tracking

- Challenge
 - Object tracking is important pipeline use case for Edge AI.
 - Identification of number of video streams with object tracking on Edge AI System is crucial.
- Solution
 - People Detection and Tracking: This pipeline uses YOLO object detection, person re-identification, and Deep SORT (Simple Online and Realtime Tracking with a Deep Association Metric) tracking, which combines:
 - Motion prediction: Kalman filter for predicting object position based on velocity and position history
 - Appearance features: Deep learning re-identification model that extracts 128-dimensional feature vectors to distinguish objects
 - Robust association: Combines IoU and cosine distance metrics to match detections with existing tracks
- Benefits
 - Identification of maximum number of video streams including object detection and object tracking

Pipeline	Description	Variants
	People Detection and Tracking: People detection and tracking pipeline using YOLO object detection, person re-identification, and Deep SORT tracking.	Available in GPU variant.




Yolo models support by ViPPET

Added support for YOLO11/YOLO26 models

- Challenge
 - Evaluate and benchmark industry-standard YOLO object detection models

Solution

- Added support for the following Ultralytics models:
 - YOLO26 s/n with Pose & Segmentation variants
 - YOLO11 s/n with Pose & Segmentation variants
- Benefits
 - Zero-setup benchmarking for popular YOLO models.



ViPPET feature structure

1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration

2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)

3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing


4 Enhanced User Experience & Visualization

- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode

In-app Performance Benchmarking with concise reporting


Streamlined benchmarking and reporting for predefined workload sets

- Challenge
 - Running consistent, repeatable end-to-end pipeline benchmarks across different HW targets and workloads required manually setting up and comparing each test
- Solution
 - ViPPET's preconfigured benchmarks can be used for comparing how sets of AI pipelines behave across device variants and stream counts. We created first benchmark suites for Retail, Metro, and Manufacturing.
 - Each suite is a collection of workloads, and each workload maps to a specific pipeline plus a set of test cases like CPU/GPU/NPU at 1, 4, 8, and 12 streams.
 - Starting a benchmark launches a job that runs the suite's test cases sequentially, tracks live progress, and prevents multiple suites from running at once. As results come in, ViPPET aggregates them into workload-level and suite-level scores, with separate performance and efficiency components plus an overall score.
 - Users can go one step deeper into individual test runs to view time-series charts for frame rate, memory, and CPU usage. Finished runs can be exported as CSV & PDF.
- Benefits
 - Users can run a repeatable, opinionated benchmark pack instead of assembling one-off performance tests by hand



The screenshot shows a detailed benchmark result for the 'Simple NVR' workload. The table displays performance metrics for different device variants and stream counts. The 'Variant' column is highlighted with a blue border.

Variant	Streams	Duration	Total FPS	Per-stream FPS	CPU	GPU	NPU	Media	Memory	Power	Status
CPU	1	0m 1s	384.94	384.94	9.8%	5.9%	59.0%	0.0%	20.7%	15.9 W	Passed →
CPU	4	0m 1s	1411.90	352.98	13.7%	0.0%	29.3%	0.0%	19.0%	771 W	Passed →
CPU	8	0m 1s	1995.85	249.48	41.2%	0.0%	0.0%	0.0%	19.3%	45.7 W	Passed →
CPU	12	0m 2s	1809.50	150.79	61.3%	0.0%	0.0%	0.0%	20.1%	48.5 W	Passed →
GPU	1	0m 0s	1157.83	1157.83	93.9%	0.0%	0.0%	0.0%	21.0%	79.9 W	Passed →
GPU	4	0m 1s	1089.32	272.33	12.6%	0.0%	0.0%	0.0%	19.2%	33.6 W	Passed →
GPU	8	0m 2s	1131.00	141.38	3.1%	0.0%	0.0%	0.0%	19.1%	9.9 W	Passed →
GPU	12	0m 3s	1115.50	92.96	3.0%	0.0%	0.0%	68.5%	19.5%	10.7 W	Passed →



Extended density benchmarking

Identify and maximize available platform headroom to safely run additional edge AI workloads

- Challenge
 - ViPPET supported 2-pipeline density testing, but lacked the control to set a specific, fixed stream count for one pipeline leaving available headroom unmeasured.
- Solution
 - Lock the exact stream count for Pipeline #1 (baseline workload), then automatically scale Pipeline #2 to find the hardware limit.
 - Extended ViPPET density benchmarking so users can fix the count of input streams for one pipeline and then automatically increase the count of input streams for the second pipeline class until the platform limit is reached.
- Benefits
 - Identify how available headroom can be used to safely deploy secondary analytics without compromising primary stream SLAs.

Density Tests
Density test finds the maximum number of streams per pipeline for the specified minimum FPS per stream

☐ Classic ☒ Mixed

Exactly two pipelines. The first pipeline runs a fixed number of streams, the second one is incremented by the benchmark algorithm until the target FPS is no longer met.

Pipeline	Variant	Streams
Age & Gender Recognition	CPU	8
Defect Detection	CPU	To be calculated

Set target FPS
30 FPS

☐ Enable latency metrics ☐ Set iteration duration


[Run density test](#)

NPU Metrics


ViPPET now supports reporting NPU utilization

- Challenge
 - ViPPET 2026.0 reported utilization only for CPU and GPU
- Solution
 - ViPPET now supports reporting NPU utilization.
 - NPU (Neural Processing Unit) is a dedicated hardware accelerator integrated directly into modern Intel processors (such as the Core Ultra series) specifically designed to handle artificial intelligence (AI) and machine learning tasks efficiently.
- Benefits
 - NPU is specifically designed to handle artificial intelligence (AI) and machine learning tasks efficiently – it is critical to be able to observe how pipelines perform on NPU


NPU Usage Over Time



NPU Power Over Time




NPU Frequency Over Time



Pipeline Latency Reporting

Added support for reporting latency metrics to show end-to-end pipeline processing time

- Challenge
 - ViPPET 2026.0 measured only throughput of DLS pipelines
- Solution
 - This release added support for reporting latency metrics to show end-to-end pipeline processing time
 - Latency metrics show end-to-end pipeline processing time — how long each video frame takes to travel from the source element to the sink element
- Benefits
 - Latency metrics help to identify bottlenecks and evaluate whether a pipeline meets real-time requirements



ViPPET feature structure

1 Data Ingestion and Customization

- Custom Pipeline Creation & Templates
- Custom Data Inputs
- Custom Model Integration

2 Pre-Built AI Workloads

- Industry-Specific Pipelines
- Advanced AI Capabilities (e.g. segmentation, pose estimation, DeepSORT)
- Broad model support (e.g. Yolo)

3 Comprehensive Performance Analysis & Benchmarking

- Holistic System Metrics
- Benchmarking and Reporting
- Latency Reporting
- Advanced Density Testing

4 Enhanced User Experience & Visualization


- Simplified & Advanced Views
- Live Pipeline Preview
- Demo Mode

Simple View



Simplified View for Pipeline Graphs

- Challenge
 - ViPPET 2025.2 offered full GStreamer graph view only
- Solution
 - This release introduces the **Simple view** for pipeline graphs in the Pipeline builder.
 - A simplified pipeline view is available alongside the advanced full GStreamer graph view.
- Benefits
 - While the advanced view can be suitable for users who understand GStreamer, the simple view makes understanding the pipelines much easy for those who are not familiar with it.

The same pipeline shown



in *Simple View* ...




and in *Advanced View*.

Live Pipeline Output

Enabling Real-time Preview of Pipeline Output During Execution

- Challenge
 - ViPPET 2025.2 supported only saving the output to the file
- Solution
 - This release introduces live pipeline output preview
 - A WebRTC-based video player is integrated into the UI, enabling real-time preview of pipeline output during execution.
 - Works for both video and camera inputs.
- Benefits
 - Instantly see what the pipeline is producing
 - No need to wait for full pipeline completion to evaluate results.
 - Better UX/DX, stronger demos and presentations.



Demo Mode

Present in Full Screen Performance of Video Pipelines

- Challenge
 - OEMs/ODMs need to showcase their hardware capabilities to customers and partners in compelling, real-time demonstrations.
 - Partners frequently ask: "How many AI pipelines can your platform handle simultaneously? What are the actual frame rates, CPU load, and GPU utilization under real workloads?"
- Solution
 - ViPPET demo mode creates an immersive full-screen experience that runs multiple AI pipelines continuously, providing live performance visualization on large displays.
 - Multiple AI pipeline combinations execute simultaneously (for example, Smart NVR + License Plate Recognition)
 - Two comprehensive testing modes: throughput test and density test
 - Real-time visualization displays FPS, CPU/GPU utilization, power consumption, and thermal performance with interactive charts

Benefits

- Provides ISVs and customers with compelling demonstrations showcasing true OEM hardware capabilities in realistic scenarios.

