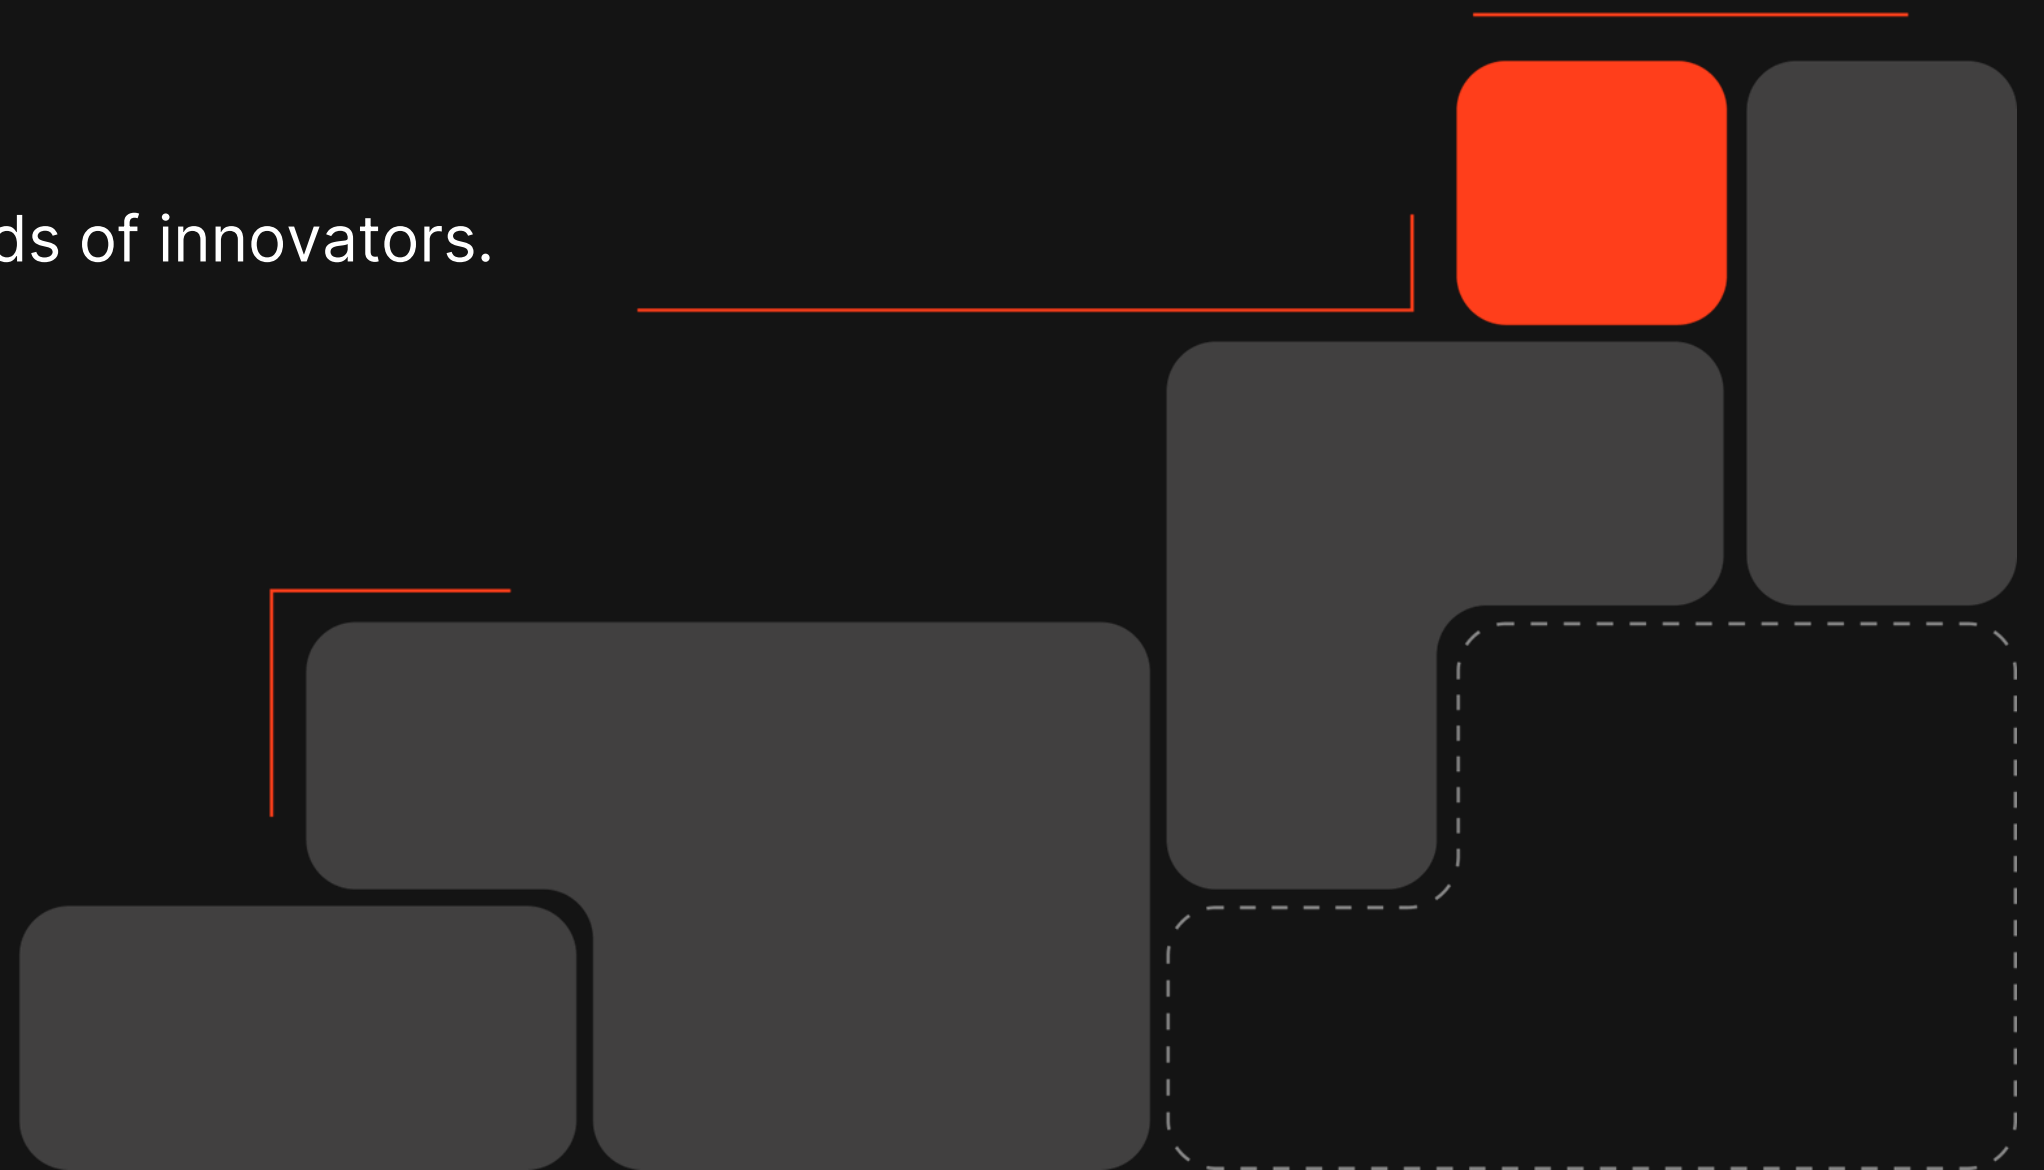


ENABLE EDGE AI PRODUCTS YOU DREAM OF

Simply bringing the power of chips into the hands of innovators.

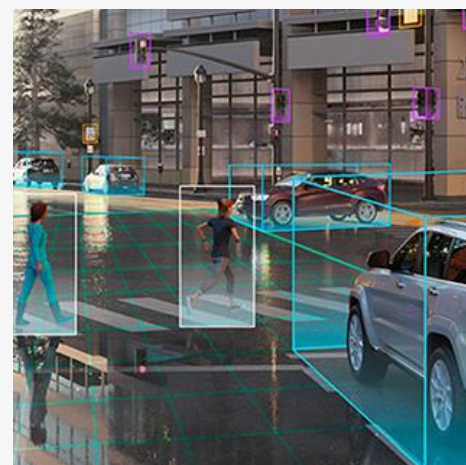
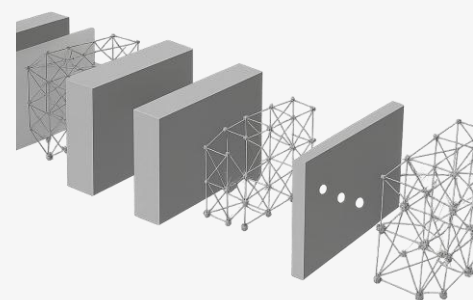


PROBLEM

BUILDING EDGE AI PRODUCTS IS STILL PAINFUL

Proof of concept

Trained ML model



3-6 months

"My model is not supported."

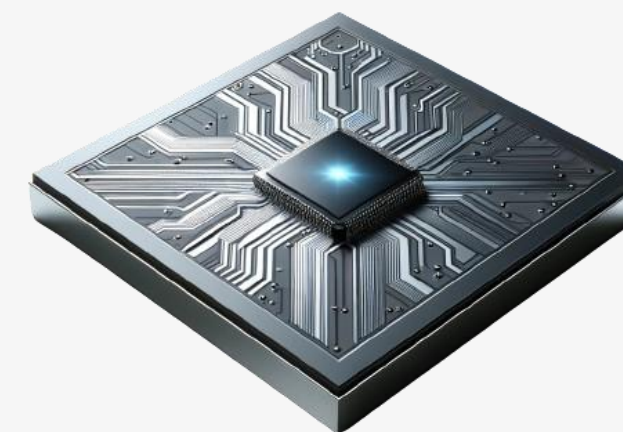
"My model is too slow."

"How does the deployment work?"



Edge AI product

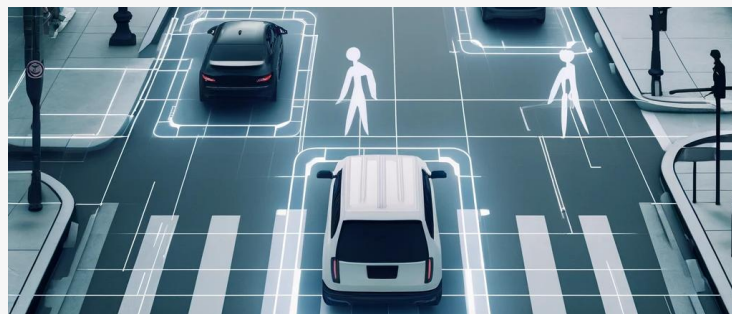
Diverse, complex, and heterogeneous chips



SOLUTION

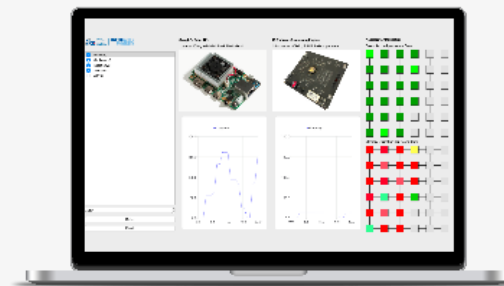
WE ENABLE CHIPS TO EFFICIENTLY RUN ANY MODEL

Any model



 **roofline**

Software development kit (SDK)



Flexible



Efficient

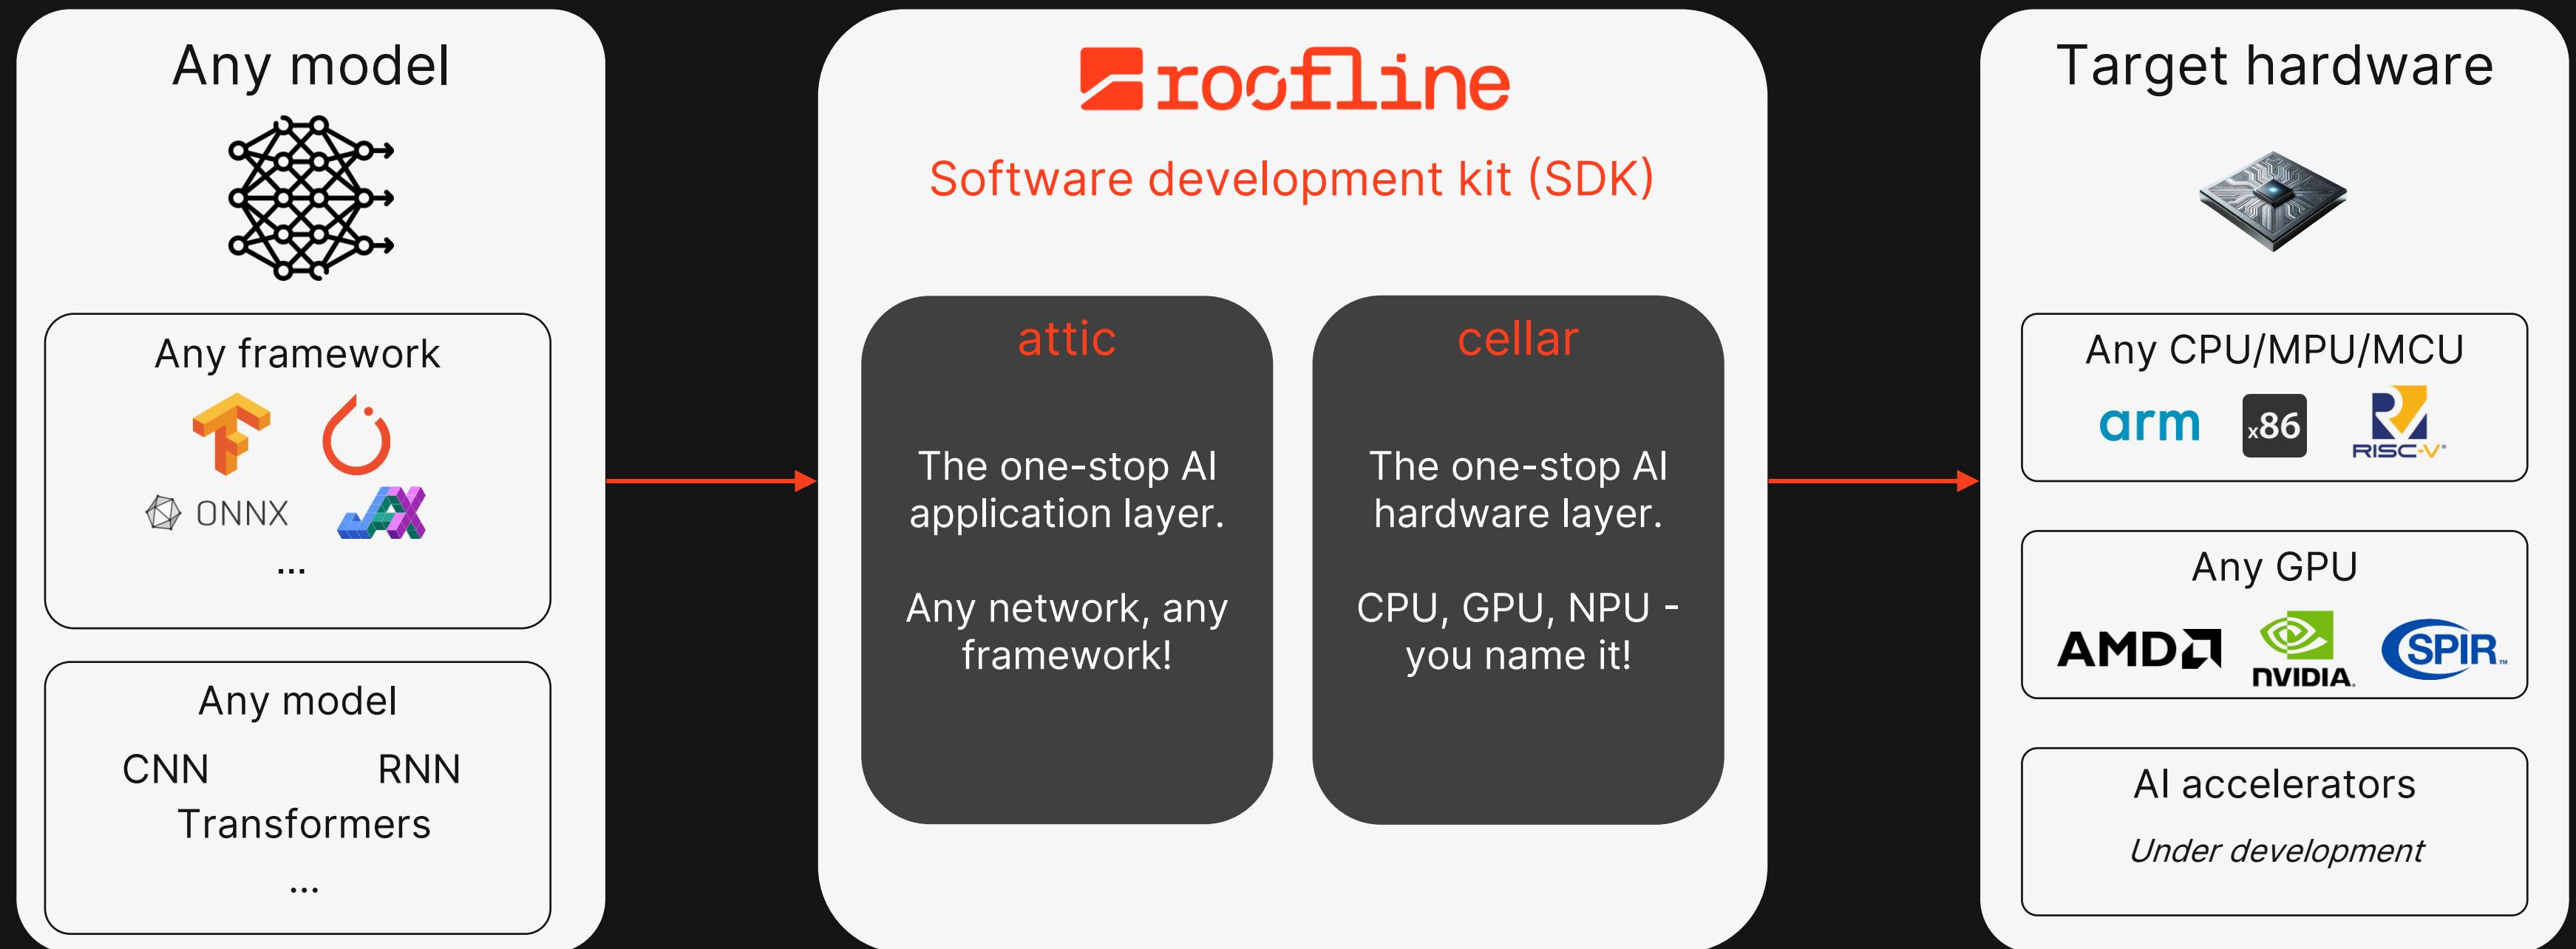


Easy

Target hardware



AN ADAPTIVE AND RETARGETABLE SDK BASED ON NEXT-GEN AI COMPILER TECHNOLOGY



WE HAVE DAY-0 SUPPORT FOR NOVEL, MULTIMODAL LLMS



- Multimodal model supporting text-to-text as well as image-to-text and text-to-image
- Not supported by llama.cpp – bug report has been open for ~2 months
- We could run the model out of the box
- We achieve ca. 8x improved latency vs TorchInductor

[Find the demo here \(link\)](#)



```
ziereis@4090 ~/r/roof-mlir (03-12-rooflm_part8_entry_and_docs)> rooflm generate janus_f16.vmbf --tokenizer-id="deepseek-ai/Janus-Pro-1B" --protocol=roofqx --prompt="My favorite color is orange because " --stream --max-new-tokens=100 2> /dev/null
Python version is above 3.10, patching the collections module.
Running warmup...
```



```
ziereis@4090 ~/r/roof-mlir (03-12-rooflm_part8_entry_and_docs)> rooflm generate "deepseek-ai/Janus-Pro-1B" --tokenizer-id="deepseek-ai/Janus-Pro-1B" --protocol=torch --prompt="My favorite color is orange because " --stream --max-new-tokens=100 2> /dev/null
Python version is above 3.10, patching the collections module.
Running warmup...
Generating text for prompt: My favorite color is orange because
My favorite color is orange because
```

Please note: janus-pro-1b f16 running on AMD Ryzen 9 7950X

ABOUT US

WE BRING THE EXPERTISE TO ENABLE YOUR AI DEPLOYMENT

Founded by AI compiler and AI system engineers from RWTH Aachen University.

Built around a core team of ex-AMD engineers, researchers and strategists.

Multi-million funding secured from public institutions and private investors.



CEO
Dr.-Ing. Moritz Joseph

joseph@roofline.ai



www.roofline.ai



CTO
Max Bartel

bartel@roofline.ai



CFO & COO
Dr. Thomas Zimmermann

zimmermann@roofline.ai



eXIST

SPRIN-D

INNOVATION
SPRINT

EXCELLENZ
**START-UP
CENTER**
NORDRHEIN-WESTFALEN

