Silicom Ltd.
Connectivity Solutions

intel ai

AI-Enabled Edge Networking and Security

Presented by

Wen Wang, Principal Software Engineering Lead at Silicom
Paul McFall, CTO & Engineering Manager, Edge Networking Division at Silicom
Kenny Fong, Solutions Architect in ECND (NEX) at Intel Corporation

Agenda

1. Collaboration & Vision

Strategic partnership between Intel and Silicom advancing Edge AI and secure networking.

2. Silicom Ibiza 1U Universal CPE

Unified edge platform integrating network acceleration, AI inference, and vision capabilities.

3. Accelerating Security – VPP Crypto Infrastructure

Performance benchmarking for high-throughput IPSec encryption and crypto acceleration.

4. AI-Driven Network Defense

Implementation of malware and malicious URL detection models for real-time edge protection.

5. Edge AI Vision in Action

Deployment of real-time computer vision inference for surveillance, retail, and industrial use.

6. Scalability & Future Outlook

Extending the Silicom–Intel edge ecosystem toward post-quantum security readiness.

Silicom Ibiza 1U Universal CPE

Universal Edge Platform

- High-speed networking, AI acceleration, and edge vision in a low power compact system.
- Intel® CPU and integrated Intel® UHD Graphics for efficient on-device AI inference.
- Scalable CPU, memory, and storage.

Security

- Hardware-assisted crypto acceleration to improve IPSEC and cryptographic workload performance.
- Hardware RoT, Secure Boot, Verified Boot, TPM.

Flexible Design

- PCIe and M.2 expansion for add-in cards and accelerators.
- Configurable LAN/WAN: Wi-Fi 6/7, 2.5GbE, optional PoE++
- Optional Lights Out Management (LOM)
- Optional 4G LTE / 5G sub-6 connectivity (5G carrier certified)

Key Use Cases

- AI-enabled network security at the edge.
- Computer vision with low-latency inference.
- High-throughput, small-footprint uCPE replacement.

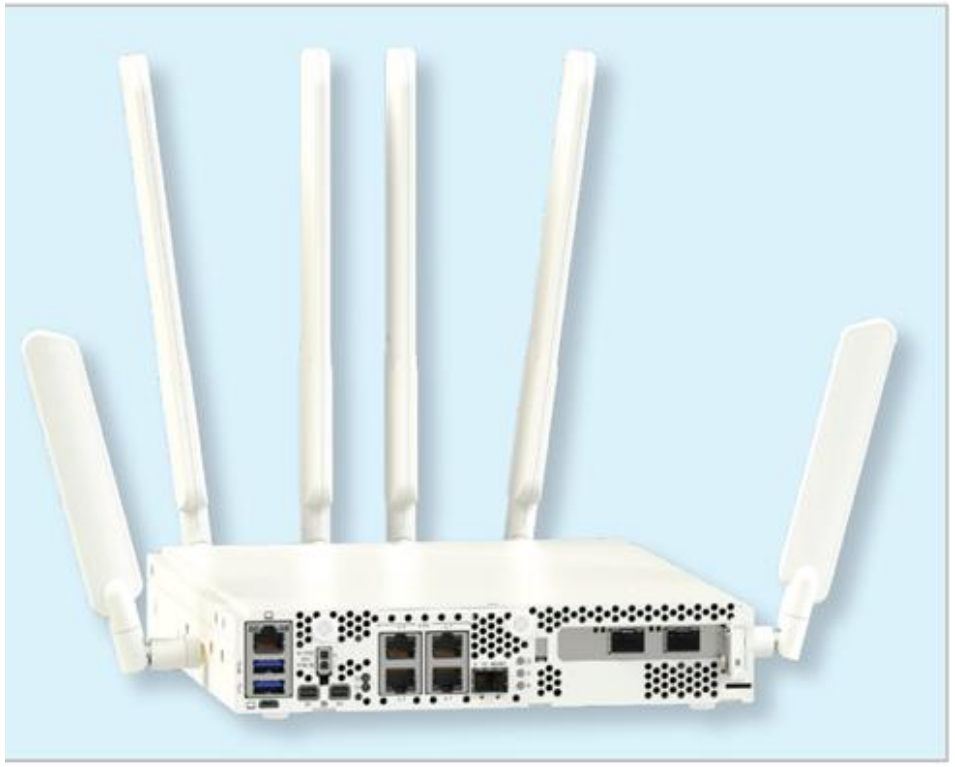

Silicom Ltd.
Connectivity Solutions

VPP Crypto Benchmarking Framework on Ibiza

- Silicom Ibiza tested with Native VPP Crypto Engine for best performance.
- vAES and vPCLMUL acceleration of AES encryption/decryption.
- Bi-directional encryption/decryption using VPP on both DUT and test server.
- DPDK pktgen provides traffic generation and throughput measurement.

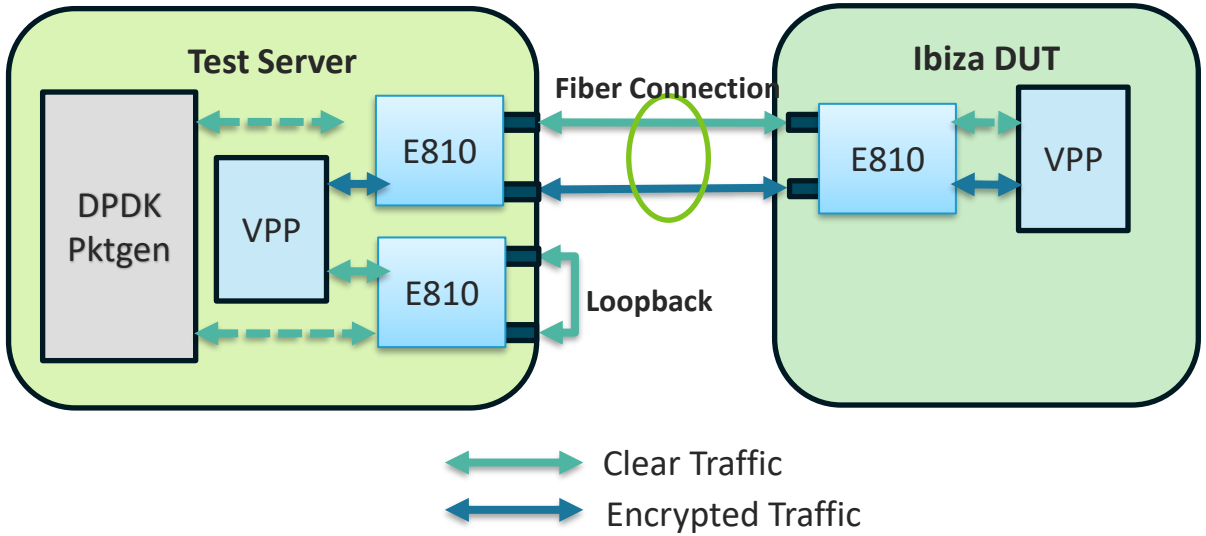

VPP-IPsec Performance on Ibiza with 2×25G NICs

- Performance measured with 1, 2, and 4 CPU cores allocated for encryption and decryption.
- Throughput scales linearly with core count, demonstrating efficiency of the native VPP crypto engine.
- Peak performance reaches 25 Gbps on 4 cores with large packet sizes.
- Results validate Silicom Ibiza as a high-throughput edge platform for secure networking workloads.

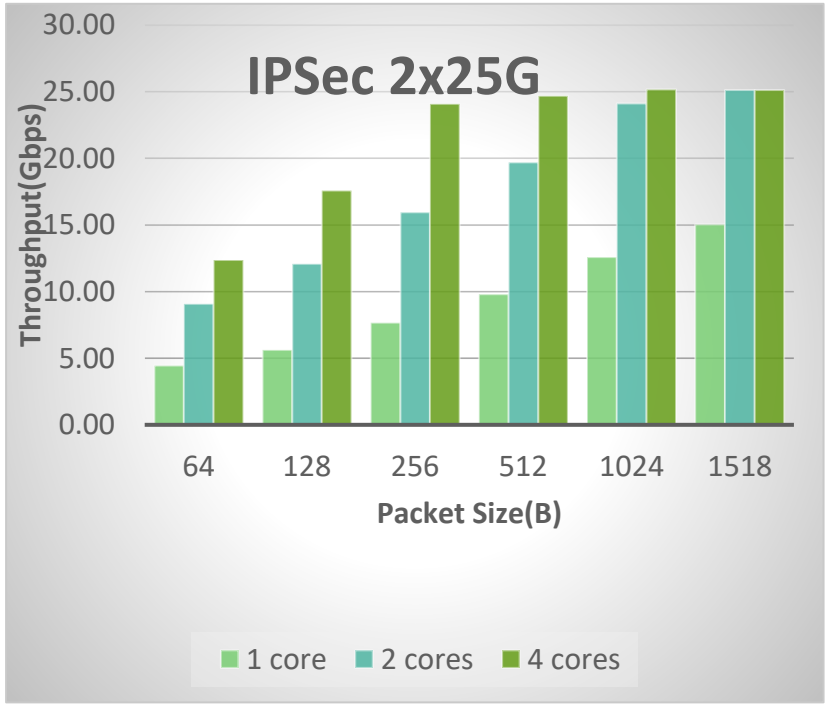

*** Throughput** is the encrypted traffic via the IPSec link between Ibiza and test server.

Performance Tuning

- **Optimize CPU Core Utilization**
  - Isolate VPP worker threads per core
  - Avoid assigning multiple devices to one thread
- **Balance Tunnels and Packet Flows:**
  - Use enough tunnels (e.g., 96 in testing) for load balance
  - Fix 5-tuple pairs in packet generator to reduce handoffs
- **Tune Descriptor Rings for Cache Efficiency:**
  - Large rings hurt cache locality — **256–512 recommended**

Silicom Ltd.
Connectivity Solutions

AI-Powered Network Security Everywhere

Industry Challenges

- AI and Automation in Cyber Attacks
- Evolving Cyber Threats
- Complex and distributed networks
- Service assurance in real-time

Customer Challenges and Opportunities

- AI democratization: Optimize AI pipelines to maximize Performance/$
- Easily integrate AI models into their solutions

Intel Response

- NetSec SW package accelerating AI solutions to market on Intel platforms
- E2E solutions under CPU/GPU
- NetSec AI Ecosystem

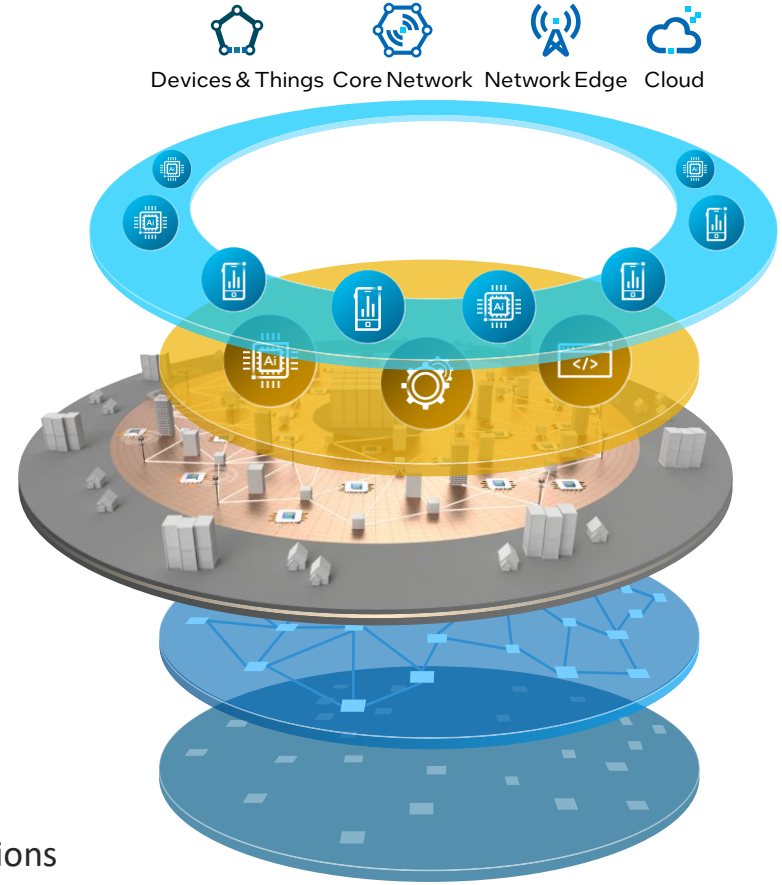

AI Everywhere

NetSec workloads need CPU, GPU, Connectivity, and SW to accelerate AI

Silicom Ltd. Connectivity Solutions

Intel Network Security AI Offering

Simplifying the AI Adoption into Network Security Workloads

- **Performance Tuning : Accuracy/Latency**
- **AI Deployment : Cloud to Edge for TCO saving**

Network Security Workloads	Intel Optimized AI models examples
Email Phishing	Intel Optimized BERT
Anomaly Detection	CNN/Rappnet (Transformer)/BERT
Malware PE *	MalConv (CNN)
Malicious URL , C2 detection	URLnet (CNN)
Malicious Packet flow	Rappnet (Transformer)
Data Loss Prevention/XDR*	BERT/LLM *
AIOps/AI Assistant	Llama, Qwen, DeepSeek

Intel Optimized Libraries, Frameworks and Toolkits

AI Tools: IPEX-LLM; Intel® Neural Compressor; Intel® Extension for TensorFlow*(ITEX); Traffic Analytics Development Kit (TADK); Intel® Extension for PyTorch* (IPEX)

Optimized Frameworks: TensorFlow, PyTorch, ONNX, OpenVINO, scikit-learn, LightGBM, XGBoost

Performance Libraries: oneAPI, oneDNN, oneDAL, oneCCL, oneMKL

Max Optimization on Intel® Processors: intel. ATOM, intel. XEON D

Accelerate with Purpose

Intel® Deep Learning Boost (Intel® DL Boost), Intel® AVX2, Intel® Advanced Vector Extensions 512 (Intel® AVX-512), Intel® Advanced Matrix Extensions (Intel® AMX)

* PE : Portable Executable , XDR : Extended detection and response, LLM (Large Language Model)

Silicom Ltd. Connectivity Solutions

NetSec AI – Overview

Comprehensive solution for NetSec enhancement with advanced AI and DL techniques.

Comprehensive AI Security Stack

- Model conversion/compression/fine-tuning/benchmarking
- Flexible API, Dockerized environment, and Intel silicon optimizations.

Supported Models & Solutions

- MalConv, URLNet – Open source models
- DNS-over-HTTPS malicious domain detection
- *Fastjoy* – high-performance flow feature extraction & analysis
- *Email Phishing Detection* – BERT-based inference integrated with Haraka mail proxy & Postfix/Dovecot

Reference Architectures

- HTTPS proxy + DNS resolver + AI inference engine
- Mail proxy + mail server + inference pipeline

Intel Hardware Acceleration

- Intel® QAT, Intel® Graphics, Intel® AVX2, and platform-wide optimizations

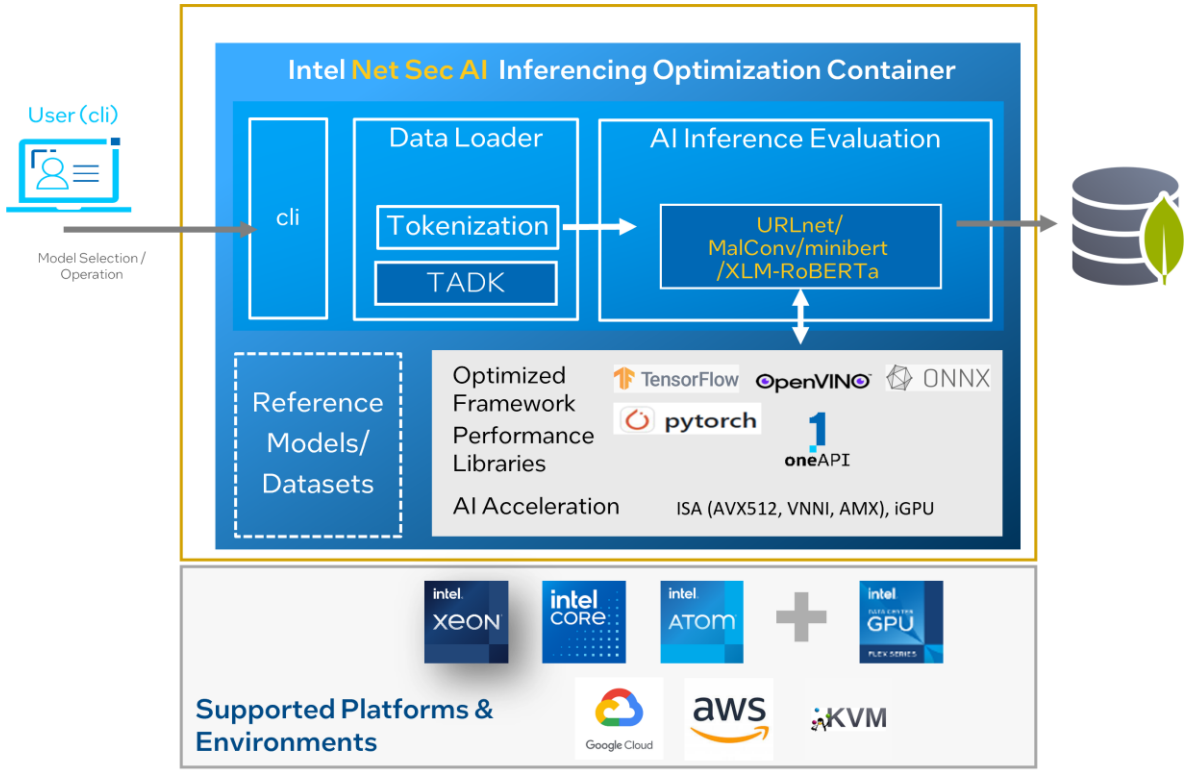

Silicom Ltd. Connectivity Solutions

Network Security – Malware Detection

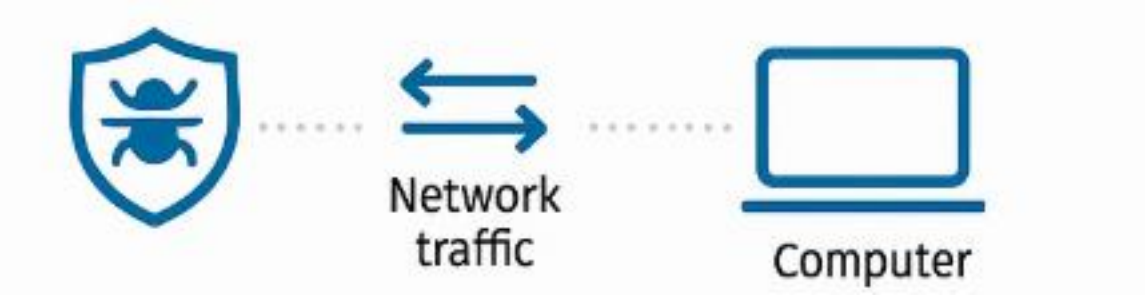

What It Does

- Scans Portable Executable (PE) files (e.g., `.exe`, `.dll`) before execution
- Uses AI models to identify malicious vs. safe files in real time
- Operates inline — no cloud sandboxing required
- Prevents zero-day threats and reduces propagation across the network

Edge Networking Use Cases

- Next-Gen Firewall (NGFW): Blocks infected files at the edge
- Intrusion Prevention Systems (IPS): Detects threats mid-transit
- Secure SD-WAN & SASE: Stops malware before branch-to-branch movement
- IoT & Edge Gateways: Protects local devices without sending data to the cloud

Network Security – Malware Detection (Intel® Turbo Boost Technology Enabled)

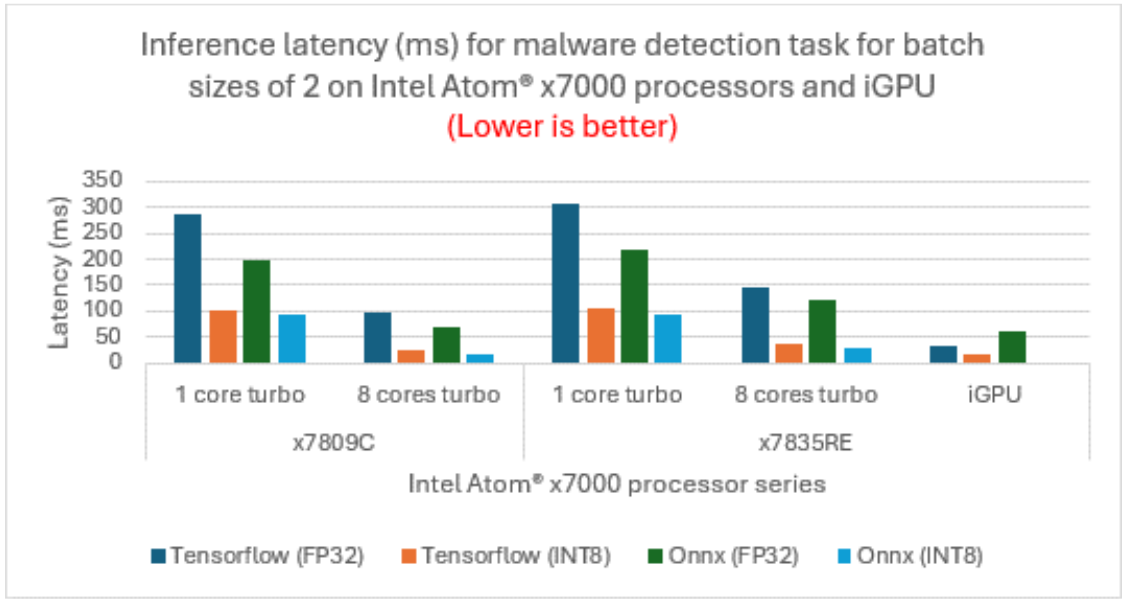

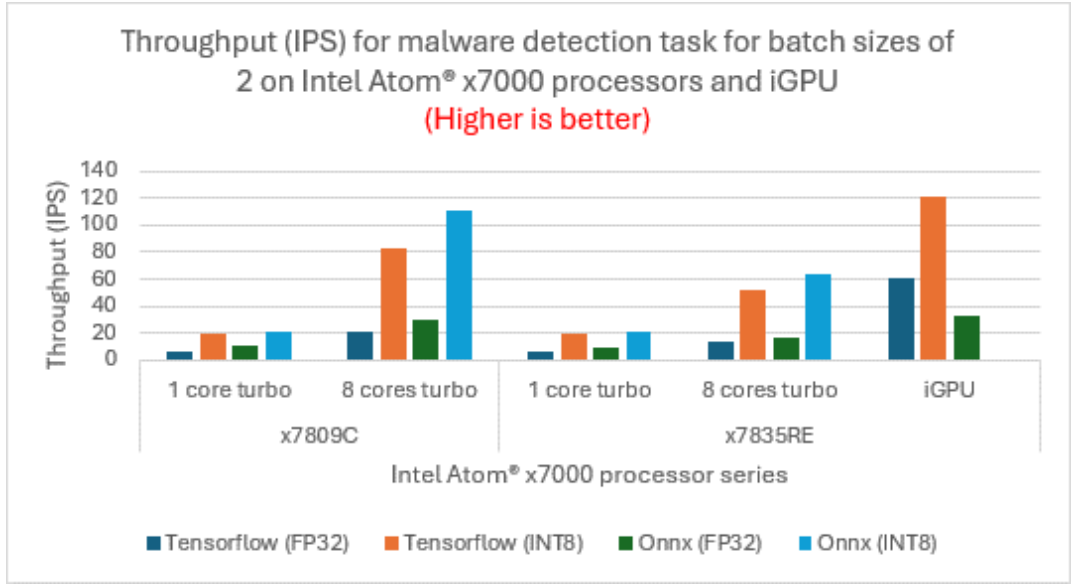

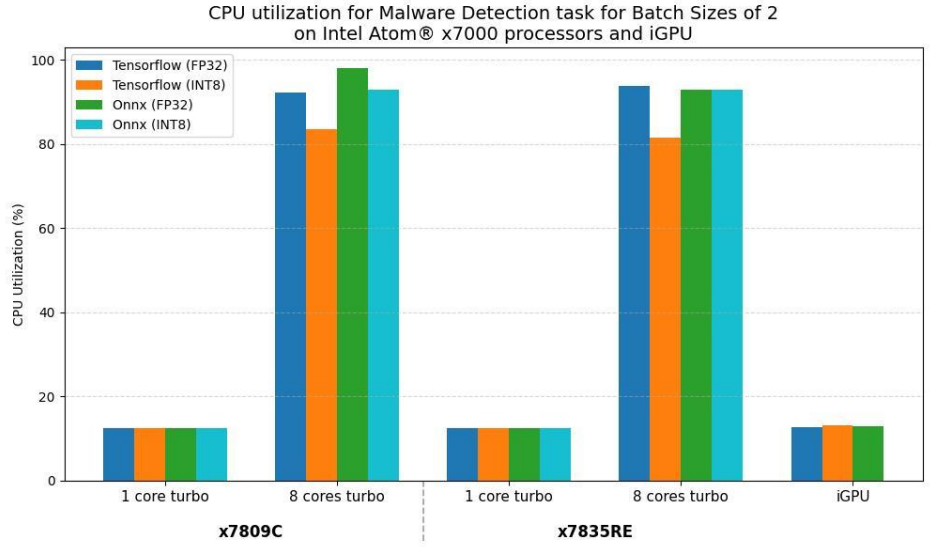

- Large scale neural network
- Quantization significantly boosts performance across both CPU and iGPU configurations
- iGPU-only execution on the Intel Atom® x7835RE processor leverages hardware parallelism and outperform 8-core CPU execution.

Silicom Ltd.
Connectivity Solutions

Network Security – URL Detection

What It Does

- **Examines web traffic for malicious URLs**
- **Uses lightweight URLNet AI models for fast real-time analysis**
- **Operates inline, no internet lookup required**
- **Detects phishing, C2 servers, and malware sites before the connection is established**

Edge Networking Use Cases

- **Secure Web Gateway (SWG)**: Filters unsafe links from employee browsing
- **Next-Gen Firewall (NGFW):** Blocks access to known malicious sites
- **Secure Access Service Edge (SASE):** Protects remote users without backhauling
- **Unified Threat Management (UTM):** Combines URL detection with other security features

Network Security – URL Detection (Intel® Turbo Boost Technology Enabled)

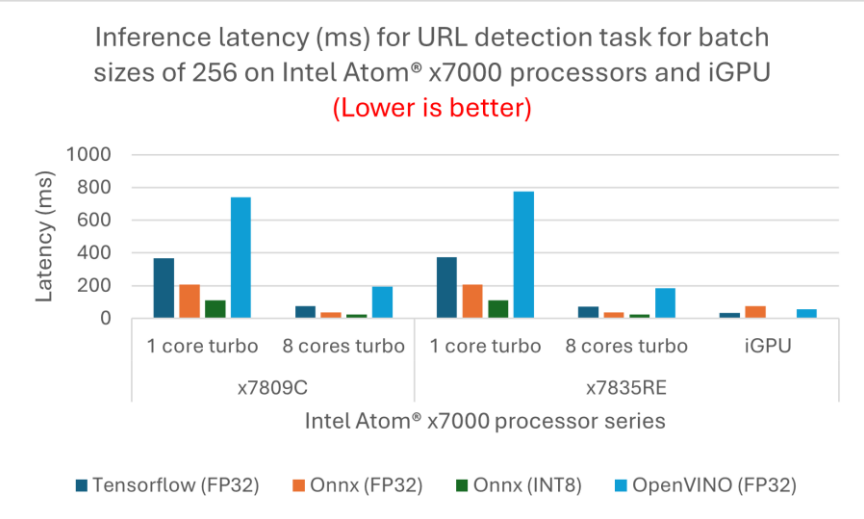

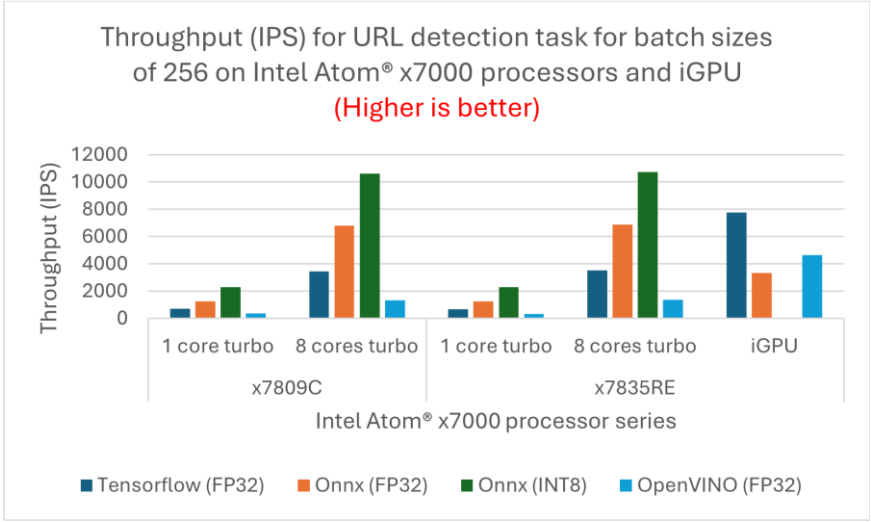

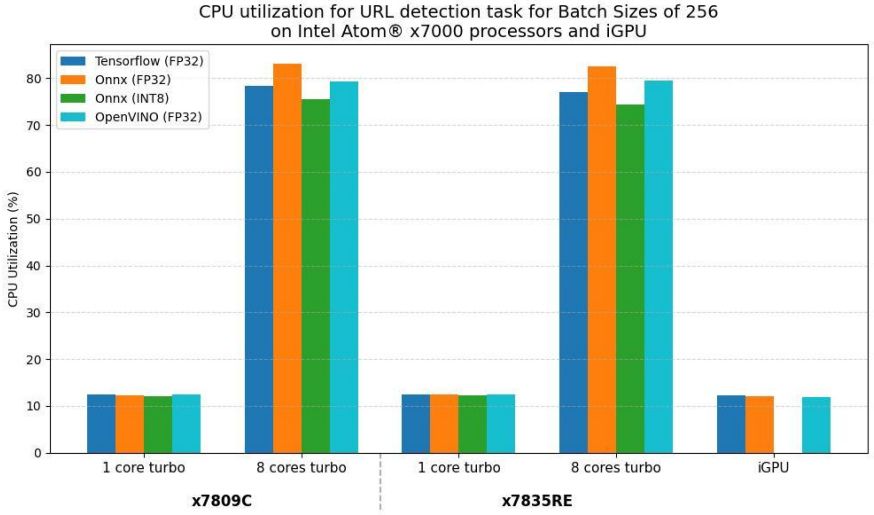

- ❑ Light weight neural network
- ❑ Need high batch size to saturate multicore-CPU or iGPU computation power
- ❑ Quantization boost the performance a lot on CPU

Silicom Ltd.
Connectivity Solutions

DNS-over-HTTPS (DoH) URL Threat Detection Demo

- **Traffic Generation:** The client sends 400 concurrent DNS-over-HTTPS (DoH) queries to simulate realistic web-browsing traffic.
- **AI-Powered Threat Detection:** The server runs each URL through an AI model (**URLNet**) to identify malicious domains, processing requests in batches of 32.
- **Performance Comparison:** Two test cases are executed
  - using **CPU inference** (2 processes)
  - using **GPU inference** (1 process)

  to evaluate inference performance across compute types.
- **Metrics:** The system measures **queries-per-second throughput**, **CPU utilization**, and **latency** to assess protection effectiveness while maintaining high performance.

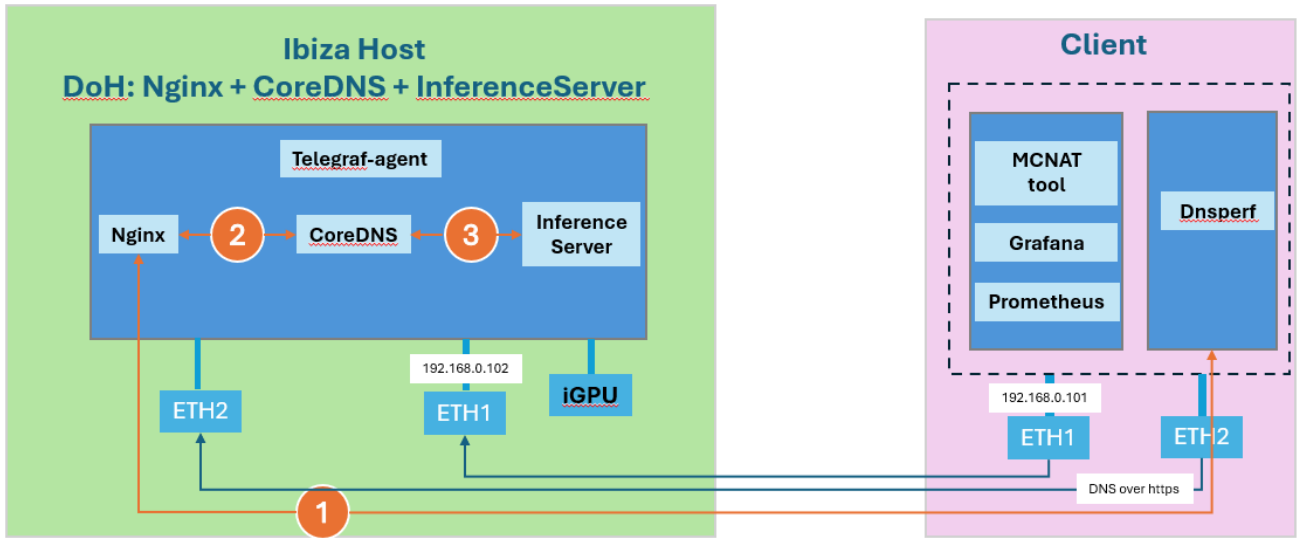

MCNAT: Multi Cloud Networking Automation Tool - from Intel

Silicom Ltd.
Connectivity Solutions

AI-Powered DNS-over-HTTPS (DoH) Security Benchmark

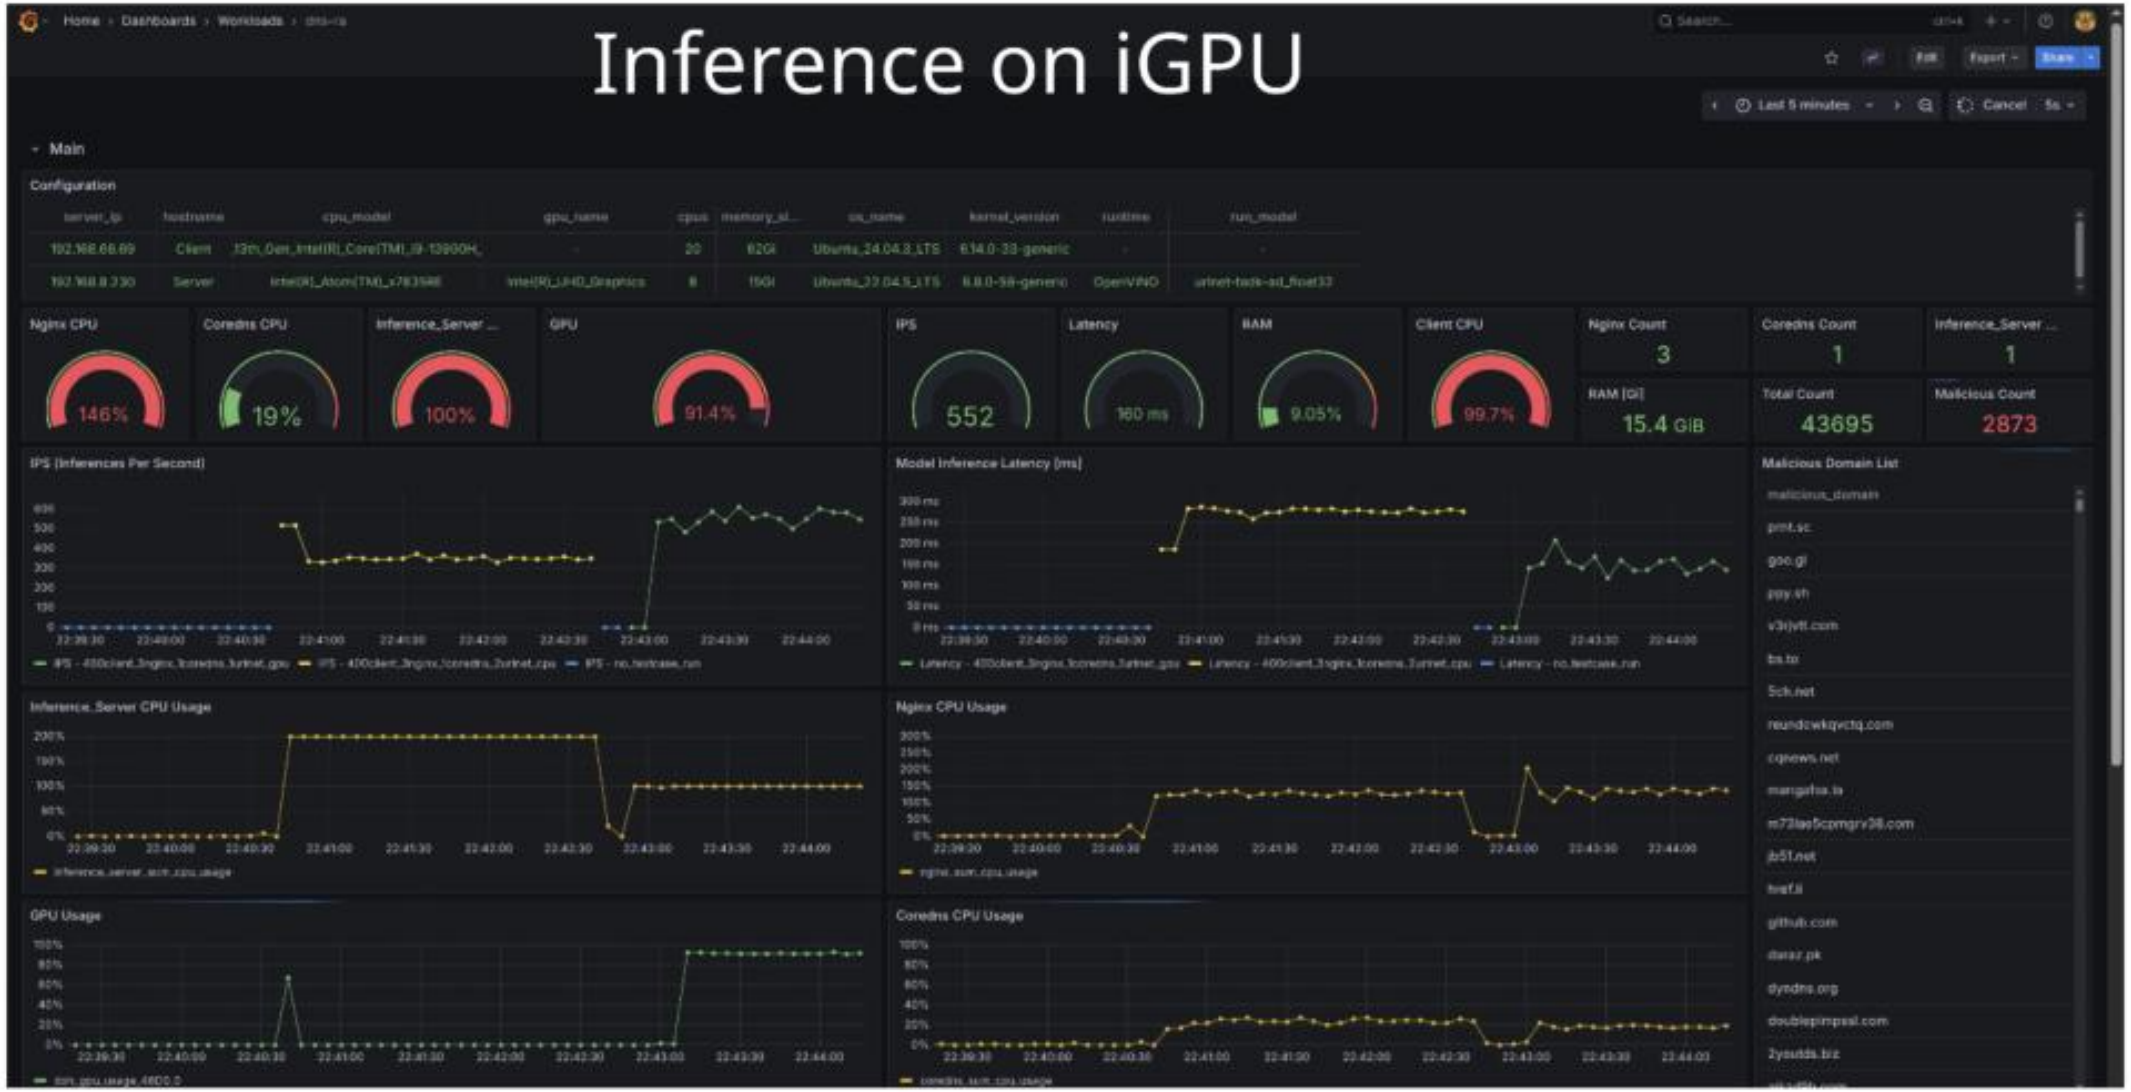

Silicom Ltd.
Connectivity Solutions

Edge AI Vision on Ibiza

- **Optimized for Modern AI Models** – Runs YOLOv11, segmentation, classification, and multi-modal inference efficiently
- **Energy-Efficient Hardware** – Significantly lower power consumption than traditional GPU servers
- **Scalable & Flexible** – Supports diverse workloads, from a few cameras to high-density multi-stream setups
- **Edge-Ready Performance** – Achieves real-time inference with smaller form factors and lower TCO
- **Broad Industry Support** – Designed to accelerate vision AI across surveillance, retail, healthcare, manufacturing, logistics, and IoT

Surveillance Application – Intel Atom® x7835RE iGPU inference

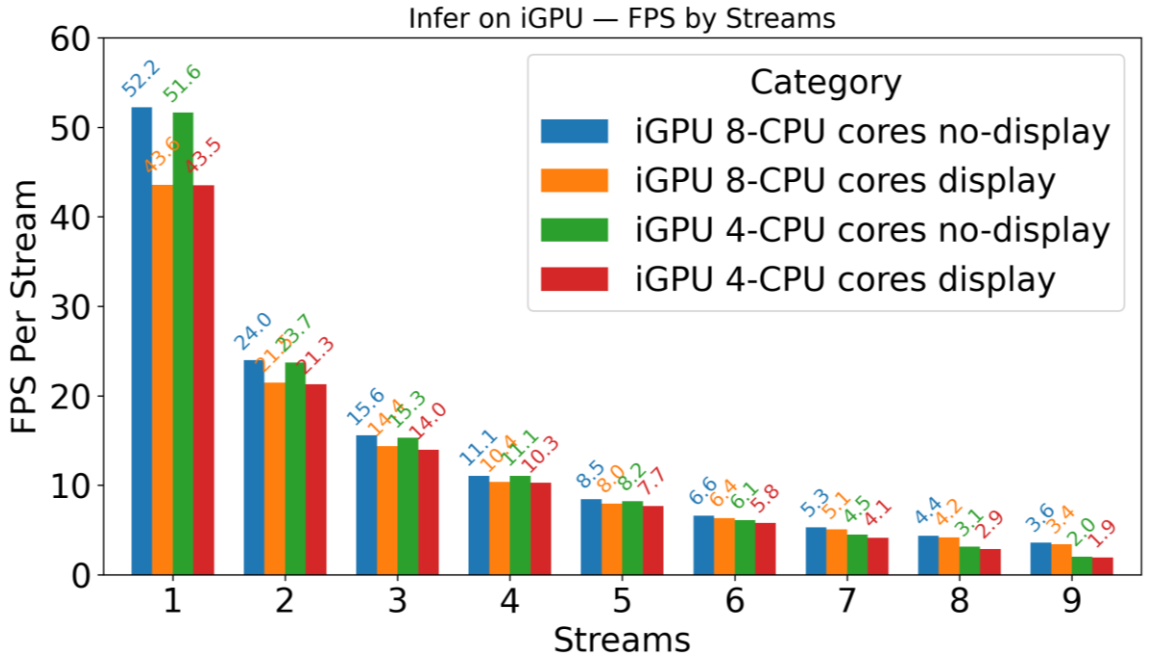

Yolov11n

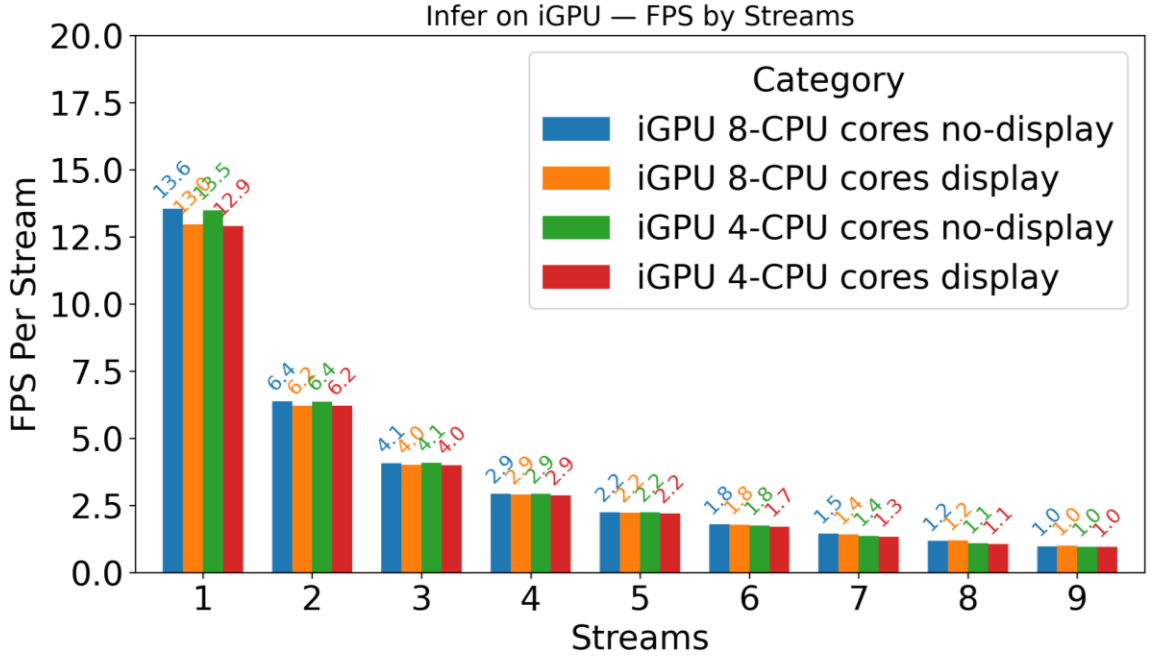

Yolov11m

- Inference on iGPU device of Intel Atom® x7835RE processor
- In general cases, FPS no-display > display, 8-core > 4-core, small model (n) > bigger model (m)
- Increase total number of streams will decrease the speed (FPS) of each stream.

Silicom Ltd.
Connectivity Solutions

Silicom Hardware Scalability

Scalable from Intel Atom® to Intel® Xeon® — A Unified Architecture for Networking, Security, and AI

Ibiza 1U uCPE Series

Processor: Intel Atom® x7000 processor series

- Compact universal CPE for SD-WAN and SD-Branch
- Up to 8 efficient cores for AI-driven security
- Built-in Intel® Xe Graphics for real-time AI inference

  Acceleration: Intel® Deep Learning Boost (VNNI), Intel® AVX2

Cadiz 1U & Desktop uCPE Series

Processor: 13th Gen Intel® Core™ processor (H/P/U)

- Enterprise-grade compute for networking and cybersecurity workloads
- Up to 14 cores / 20 threads for scalable workloads
- Built-in Intel® Iris Xe Graphics for real-time AI inference

  Acceleration: Intel® Deep Learning Boost (VNNI), Intel® AVX2

Marbella & Seville 1U Networking Appliance Series

Processor: Intel® Xeon® D-1700 D-1800, D-2700, D-2800

- Server-class performance and AI-accelerated security
- Universal CPE with up to 22 Intel Xeon cores / 44 threads
- PCIe x16 Gen4 Expansion slots for additional AI acceleration

  Acceleration: Intel® Deep Learning Boost, Intel® AVX-512, Intel® QAT

Tenerife Networking Appliance Series

Processor: Intel® Xeon® 6 SoC HCC

- High-bandwidth server-class network appliance for AI and cybersecurity at the edge
- Up to 42 Intel Xeon cores / 84 threads
- PCIe x16 Gen5 Expansion slots for additional AI acceleration

  Acceleration: Intel® AMX, Intel® AVX-512, Intel® QAT

Silicom Ltd.
Connectivity Solutions

PQC Standards and Timeline

- **Definition:** **Post-quantum cryptography** (**PQC**), sometimes referred to as **quantum-proof**, **quantum-safe**, or **quantum-resistant cryptography**, is the development of cryptographic algorithms (usually public-key algorithms) that are thought to be secure against a cryptanalytic attack by a crypto relevant quantum computers (CRQC).

- **Goal:** to develop **cryptographic systems that are secure against both quantum and classical computers** and can interoperate with existing communications protocols and networks.

NIST published PQC algorithms in 2024 and recommends everyone to **transition to PQC by 2035**.

- FIPS 203 (ML-KEM)
- FIPS 204 (ML-DSA)
- FIPS 205 (SLH-DSA)

CNSA 2.0 approved list of PQC algorithms

Digital Signatures for firmware signing	Public-key algorithms	Symmetric-key algorithms
Leighton-Micali Signature (LMS)	ML-KEM for key establishment	AES – 256 bit keys
Xtended Merkle Signature Scheme (XMSS)	ML-DSA for digital signatures	

And SHA-384 or SHA-512 for hashing as part of hardware integrity

National Security Agency

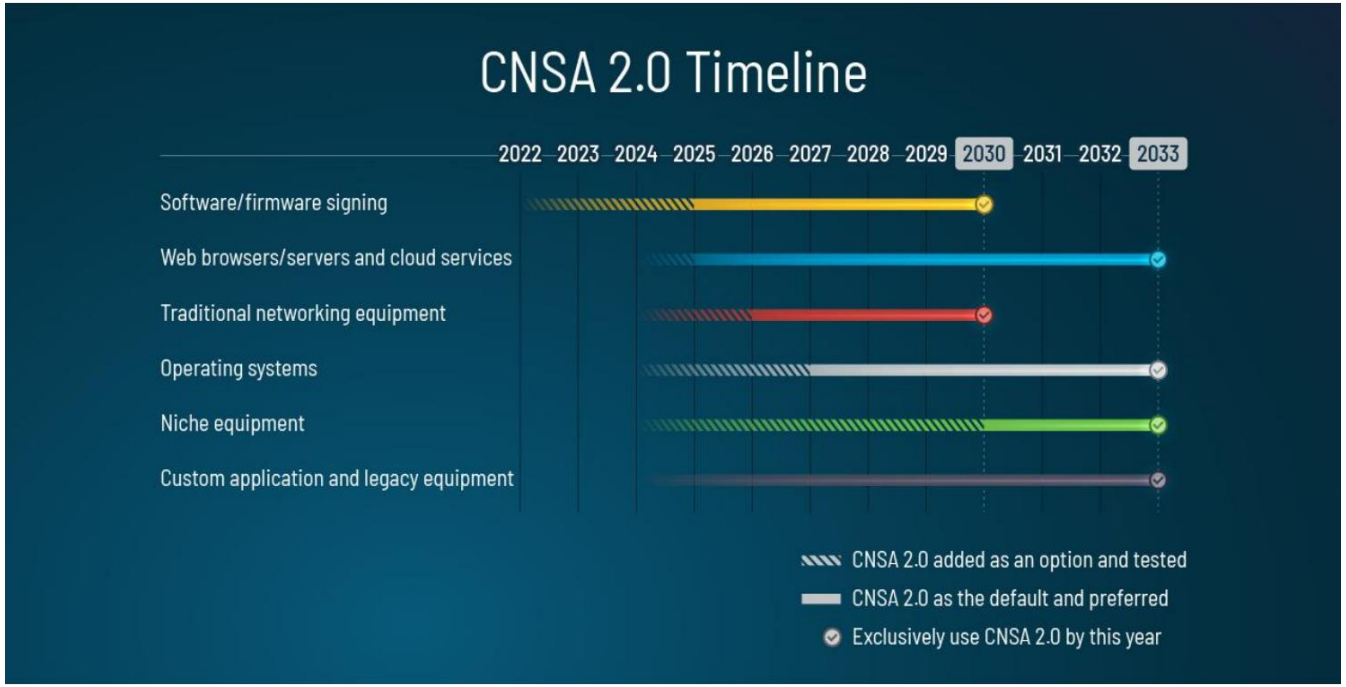

source: *NSA Cybersecurity Advisory, Announcing the Commercial National Security Algorithm Suite 2.0*

Next Steps: Bringing AI to Edge Networking Together

- **Explore Silicom's Edge Platforms**
  Evaluate Ibiza, Cadiz, and Intel® Xeon® processor-based appliances for your AI, networking, and security workloads.
- **Experiment with** NetSec software and SDK from Intel
  Integrate optimized AI models into NetSec applications using Intel-accelerated frameworks.
- **Collaborate on Performance Validation**
  Partner with Silicom and Intel engineers to tune, benchmark, and scale AI-enabled edge solutions.
- **Contribute to the Ecosystem**
  Share feedback, workloads, and deployment insights to shape next-generation edge architectures.
- **Contact Us**
  Reach out to your Silicom (contact@silicom.info) or Intel representative to discuss proof-of-concept or joint testing.

Backup Slides

VPP Crypto Infrastructure and Engines

The VPP crypto infrastructure is a crypto framework that supports different crypto engines working as plugins for high performance symmetric crypto operations. At the time of this writing, there are three crypto engines:

- Native engine: The crypto engine that is specifically designed for VPP that achieves the fastest crypto processing efficiency but with limited algorithms supported. vAES and vPCLMUL acceleration of AES encryption/decryption are automatically enabled if the application is running on the latest x86 architecture CPUs.
- IPSec-MB engine: Integration layer to Intel® Multi-Buffer Crypto for IPSec library with extended crypto algorithm support list but slightly less performance compared to the native engine. vAES and vPCLMUL acceleration of AES encryption/decryption are automatically enabled if the application is running on the latest x86 architecture CPUs.
- OpenSSL engine: The shim-layer to OpenSSL library with the most comprehensive crypto algorithm support list but is least performant.

In our test, the Native engine is used for best performance.

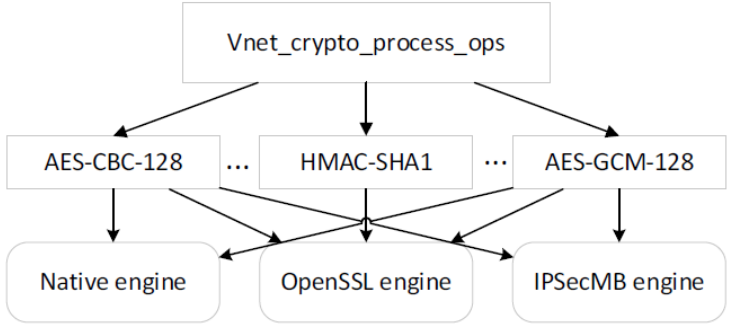

Silicom to redraw this diagram

VPP Configurations

Platform	Silicom Ibiza
CPU	Intel Atom® x7809C, 8C @ 2.4 GHz, 1 Processor, 8 Cores Intel® SSE4.1, Intel® SSE4.2, Intel® AVX2
Memory	DDR5, 4Gx4, 4800MT/s
NIC	4x2.5Gb Intel® Ethernet Controller I226-V (rev 04) 2x25Gb Intel® Ethernet Controller E810-XXVAM2 for SFP (rev 02)
BIOS	Silicom Slim Bootloader with Intel® Turbo Boost Technology enabled
OS	Ubuntu 24.04.1 LTS
Kernel	6.8.0-51-generic
Boot Settings	pcie_aspm=on intel_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=4 isolcpus=1-7 rcu_nocbs=1-7
VPP version	23.10
DPDK version	22.11.1
Pktgen version	23.03.0

AI Software Setup Configuration

Platform	Silicom Ibiza
CPU	Intel Atom® x7835RE, 8C @ 2.4 GHz, 1 Processor, 8 Cores Intel® SSE4.1, Intel® SSE4.2, Intel® AVX2
Memory	DDR5, 4Gx4, 4800MT/s
NIC	4x2.5Gb Intel® Ethernet Controller I226-V (rev 04) 2x25Gb Intel® Ethernet Controller E810-XXVAM2 for SFP (rev 02)
BIOS	Silicom Slim Bootloader with Intel® Turbo Boost Technology enabled
OS	Ubuntu 22.04
Kernel	6.8
Boot Settings	pcie_aspm=on intel_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=4 isolcpus=1-7 rcu_nocbs=1-7
VPP version	23.10
DPDK version	22.11.1
Pktgen version	23.03.0

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.