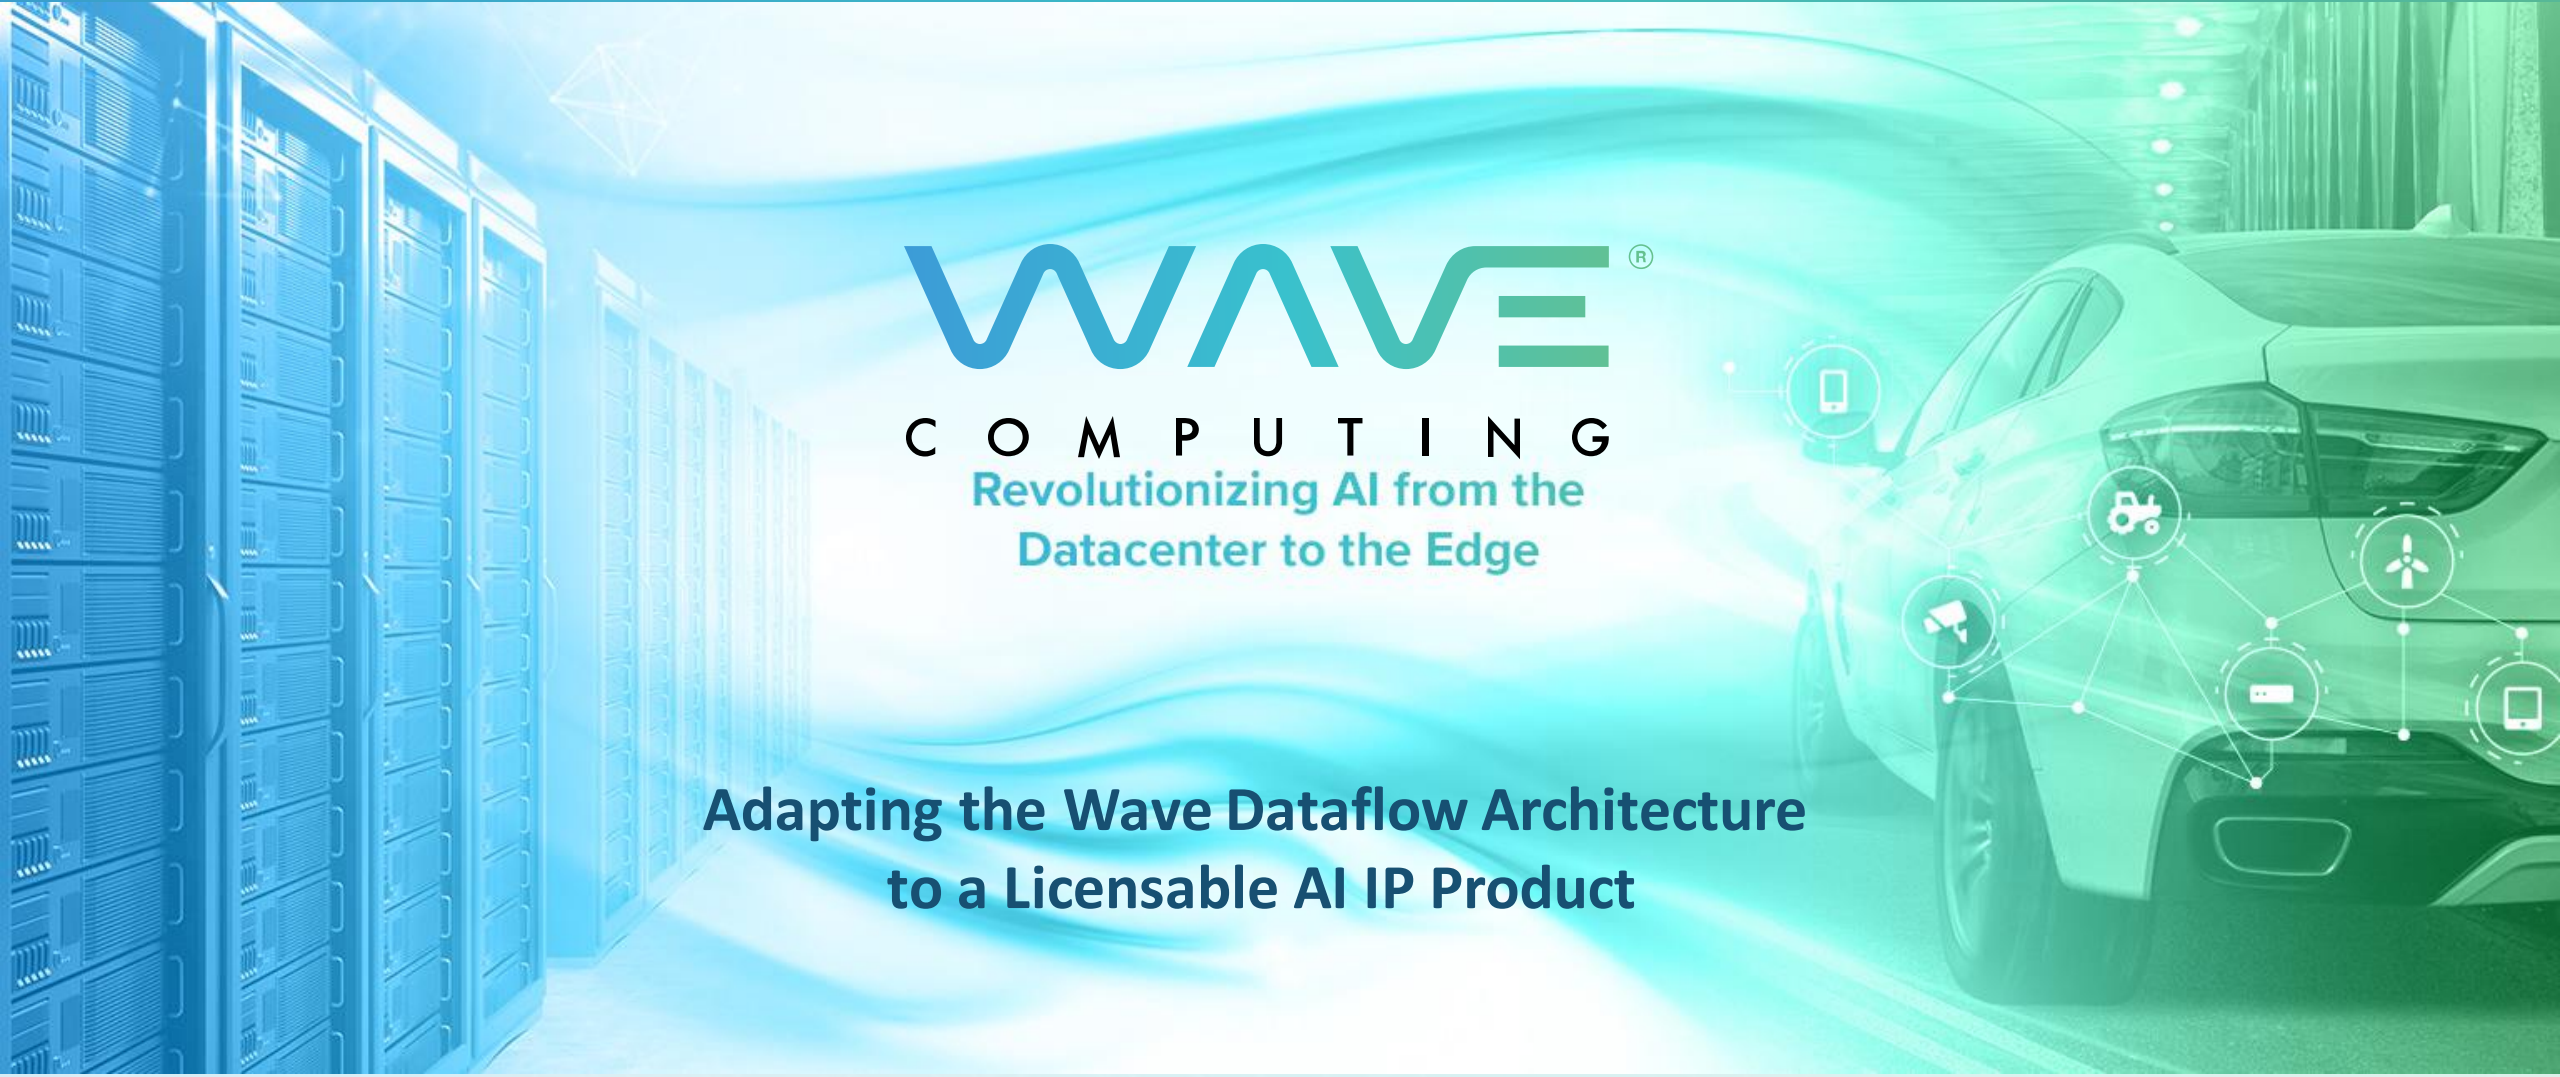

WAVE® COMPUTING

Revolutionizing AI from the Datacenter to the Edge

Adapting the Wave Dataflow Architecture to a Licensable AI IP Product

Presented by **Yuri Panchul**, MIPS Open Technical Lead

On SKOLKOVO Robotics & AI Conference. April 15-16, 2019

www.wavecomp.ai

Wave + MIPS: A Powerful History of Innovation

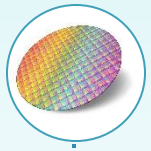

2010

Wave founded by Dado Banatao as Wave Semiconductor, with a vision of ushering in a new era of AI computing

2012

Developed Coarse Grain Reconfigurable Array (CGRA) semiconductor architecture

2014

Delivered 11GHz test chip at 28nm

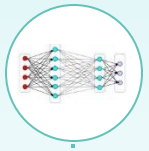

2016

Announced Derek Meyer as CEO

Launched Early Access Program to enable data scientists to experiment with neural networks

2018

Wave acquires MIPS to deliver on its vision for of revolutionizing AI from the datacenter to the edge

Announced partnership with Broadcom and Samsung to develop next-gen AI chip on 7nm node

Launched MIPS Open initiative

Closed Series D Round of funding at $86M, bringing total investment to $115M+

Expanded global footprint with offices in Beijing and Manila

2010 2011 2012 2013 2014 2015 2016 2017 2018 2019

2011

Wave develops dataflow-based technology, providing higher performance and scalability for AI applications

MIPS introduces first Android-MIPS based Set top box at CES

Google selects MIPS architecture for use in its Android 3.0, "Honeycomb" mobile device

2013

Wave expanded team to include architecture, silicon and software expertise

MIPS Technologies is sold to Imagination Technologies.

2015

Renamed the company to Wave Computing to better reflect focus on accelerating AI with dataflow-based solutions

MIPS Automotive MCUs & ADAS engagements with MobilEye, Microchip & MStar

MICROCHIP

MOBILEYE

2017

MIPS business is sold by Imagination Technologies to Tallwood Venture Capital as Tallwood MIPS Inc. for $65M

MediaTek selects MIPS for LTE modems

MEDIATEK

TALLWOOD Venture Capital

2019

Created MIPS Open Advisory Board

MIPS Open

MIPS Powering 80% of ADAS-enabled automobiles

Wave Joins Amazon, Facebook, Google and Samsung to Support Advanced AI Research at UC Berkeley

Datacenter-to-Edge Needs Systems, Silicon and IP

Compute Systems

AI Processors

AI Licensable IP

TritonAI™

Cloud

Enterprise

Semiconductor Partners

AI Applications Software

- Autonomous Driving
- NLP
- Retail
- Smart Cities

What is Driving AI to the Edge?

AI was born in Datacenter

Revolutionizing AI from the Datacenter to the Edge

Market Drivers

- Networking
- Enterprise
- Mobile
- Industrial
- Autonomous
- IOT

AI Use Cases

- Privacy
- Security
- Isolated
- Low latency

Cost

- Bandwidth
- Storage
- Compute

Scalable AI Platform: Triton*AI*™ 64

Key Benefits:

- Highly Scalable to address broad AI use cases
- Supports **Inference** and **Training**
- High flexibility to support all AI algorithms
- High efficiency for key AI CNN algorithms
- Configurable to support AI use cases
- Mature Software Platform support

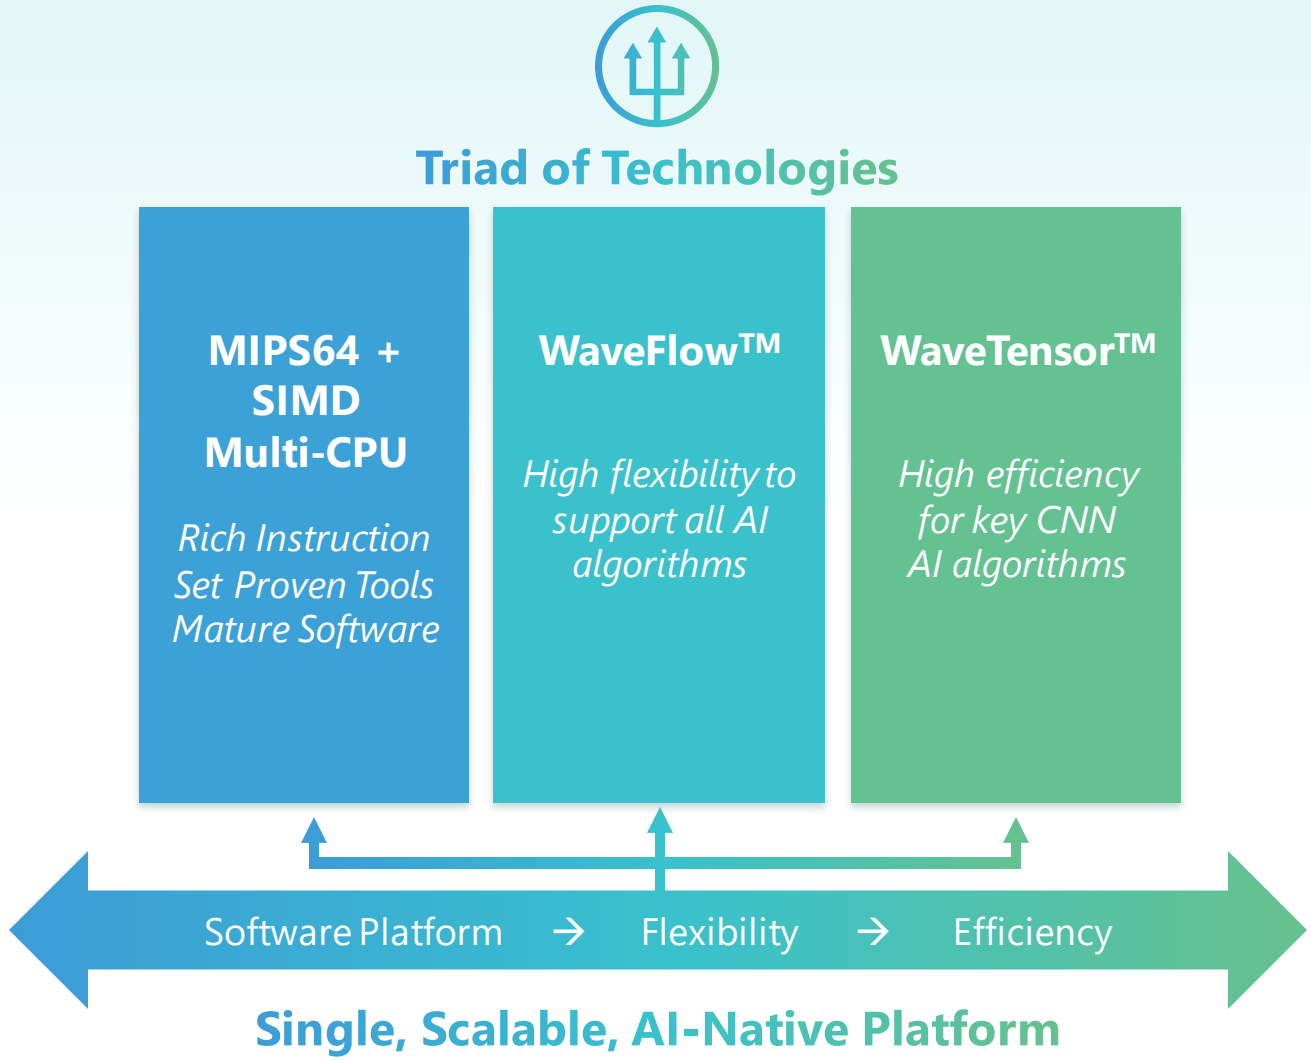

Data Rates and New AI Networks Requirements

Data Rates & Network Complexity Drive Data Compute Performance Demands

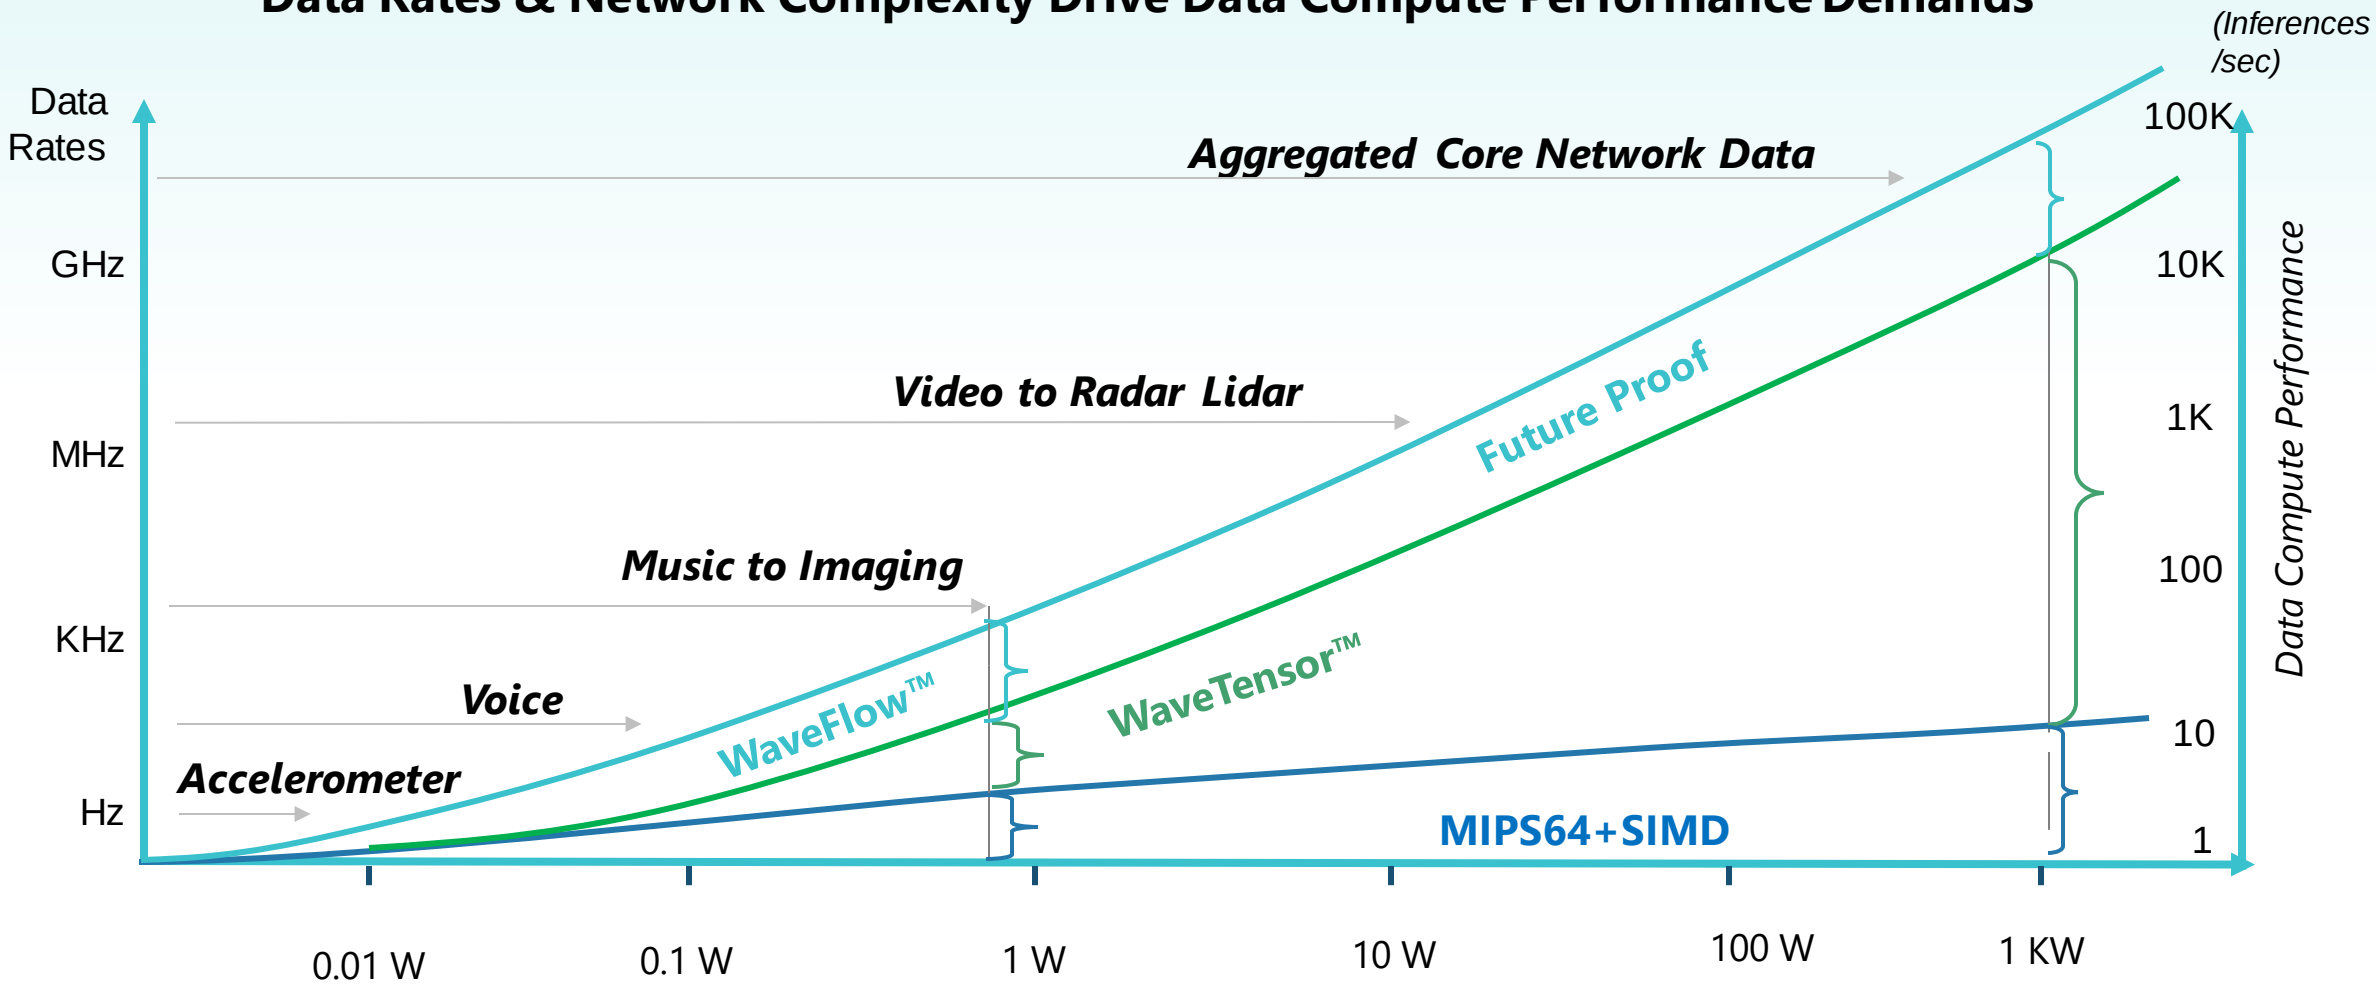

These curves represent the conceptional combination of these technologies, not actual independent performance.

WaveTensor ™ Configurable Architecture

Configurable Architecture for Tensor Processing

- Configurable MACs, Accumulation and Array Size
- Overlap of Communication & Computation
- Compatible datatypes with WaveFlow™ Core
- Supports int8 for inferencing
- Roadmap to bfloat16, fp32 for training

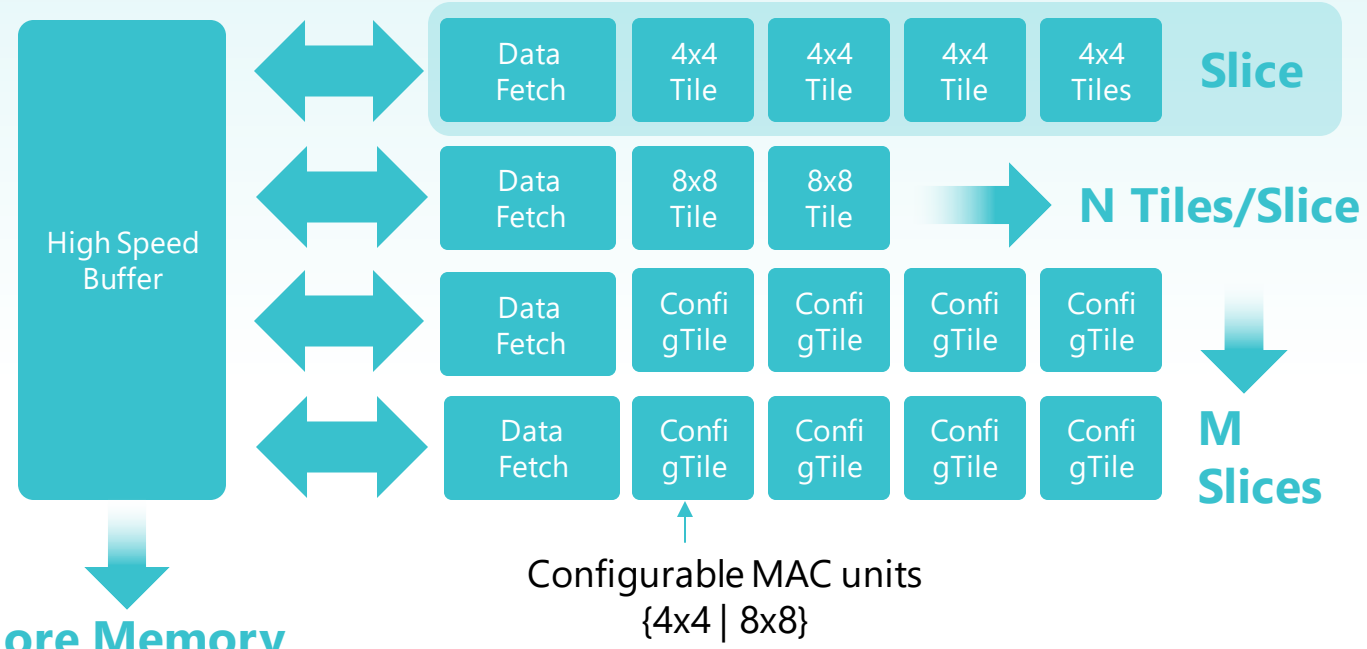

ResNet-50	Inferences/Sec*
Compute Density	~1K/mm2
Compute Efficiency	~500/watt

- Core, Int8, 7nm Fin-FET nominal process/Vdd/temp
- Batch=1, std model w/o pruning, performance and power vary with array size/configuration

MAX TOPs	TOPs/Watt	TOPs/mm²
1024	8.3	10.1

- Core, Int8, 8x8 tile config, 7nm Fin-FET nominal process/Vdd/temp

WaveFlow™ Reconfigurable Architecture

- Configurable IMEM and DMEM Sizes
- Overlap of communication & Computation
- Compatible datatypes with WaveTensor™
- Integer (Int8, Int16, Int32) for inference
- Roadmap (bfloat16, fp32) for training

- Wide range of scalable solutions (2-1K tiles)
- Future Proof all AI algorithms
- Flexible 2 dimensional tiling implementation
- Reconfigurable for dynamic networks
- Concurrent Network execution
- Supports signal and vision processing

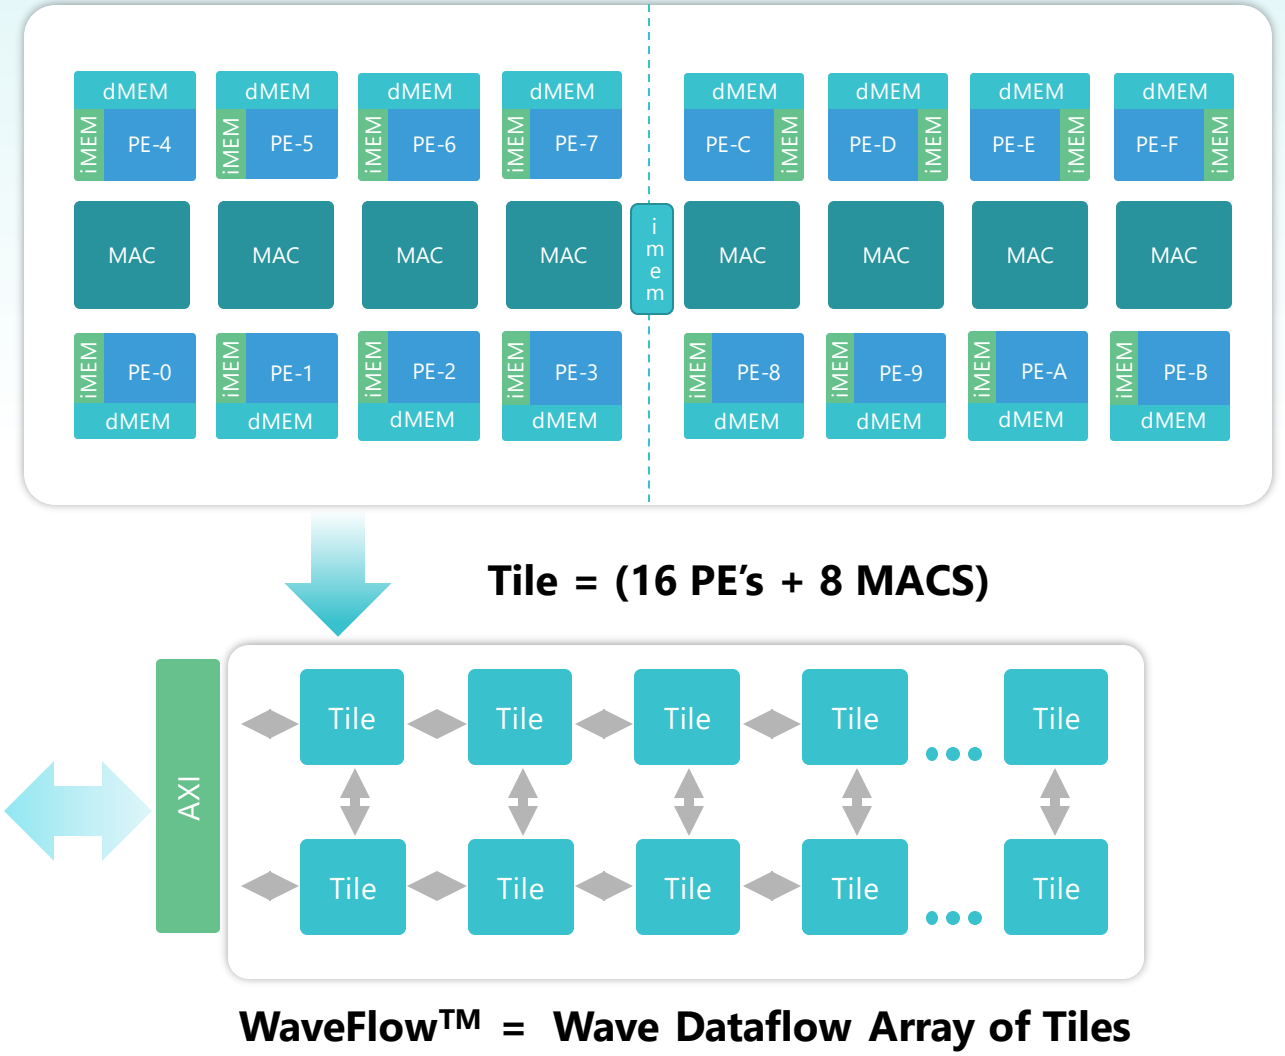

Future Proofing in an Rapidly Evolving Industry

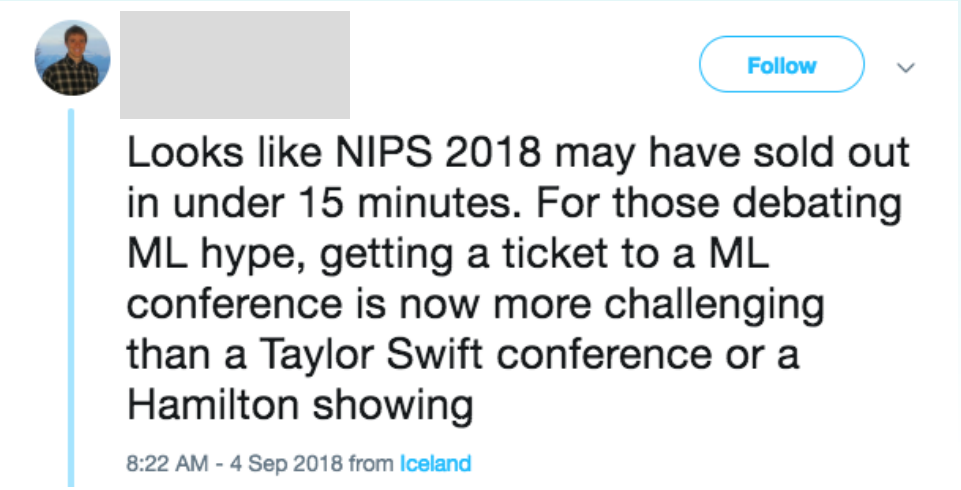

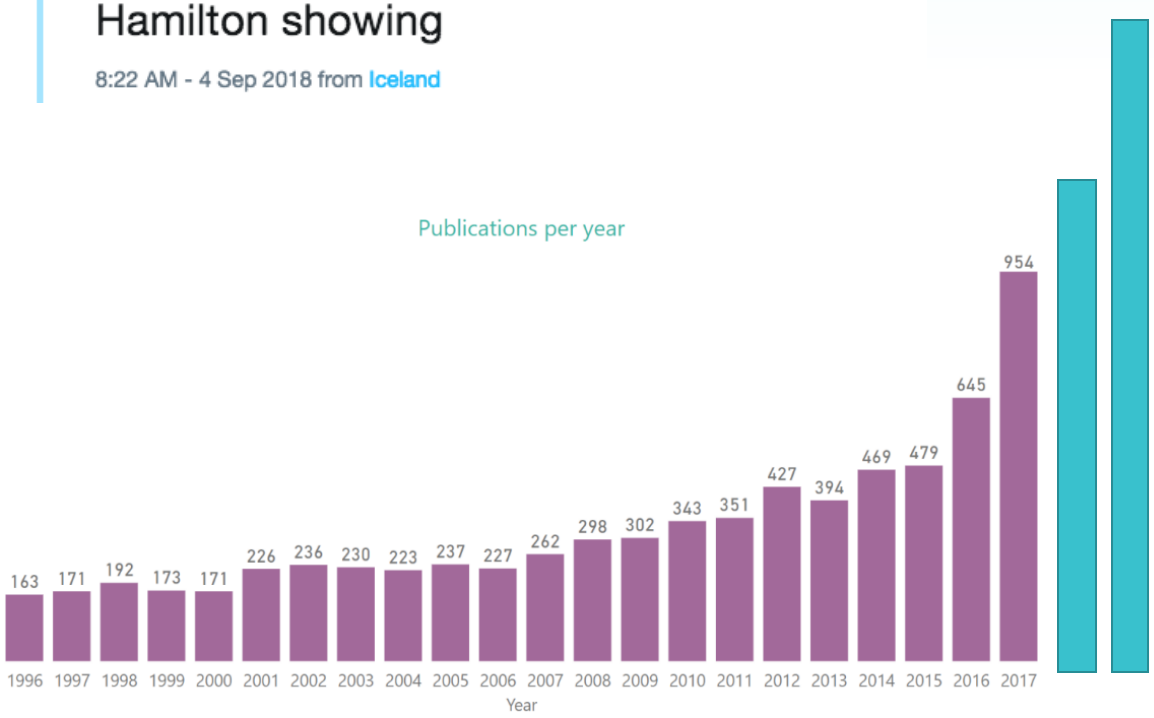

What is the likelihood that your DNN accelerator will run all these "yet to be invented" networks?

Wave's TritonAI™ 64 platform combines a reconfigurable processor with an efficient neural network accelerator.

Offers customers peace of mind and investment protection

Future-proof your Silicon

CNN Layers

- Sparse Matrix-Vector Processing
- Stochastic pooling
- Median pooling (illumination estimation & color correction)

Activation functions

- Leaky rectified linear unit (Leaky ReLU) (used in Yolo3)
- Parametric rectified linear unit (PReLU)
- Randomized leaky rectified linear unit (RReLU)

Custom Operators (e.g.)

- Novel Loss Function
- New Softmax Implementation
- Image resize nearest neighbor

Data Preprocessing

- Scaling
- Aspect Ratio adjustment
- Normalizing

Other Functions

- Compression/Decompression
- Encryption/Decryption
- Sorting

MIPS-64 Processor

MIPS-64:

- MIPS64r6 ISA
  - 128-bit SIMD/FPU for int/SP/DP ops
  - Virtualization extensions
- Superscalar 9-stage pipeline w/SMT
- Caches (32KB-64KB), DSPRAM (0-64KB)
- Advanced branch predict and MMU

Multi-Processor Cluster:

- 1-6 cores
- Integrated L2 cache (0-8MB, opt ECC)
- Power mgmt. (F/V gating, per CPU)
- Interrupt control with virtualization
- 256b native AXI4 or ACE interface

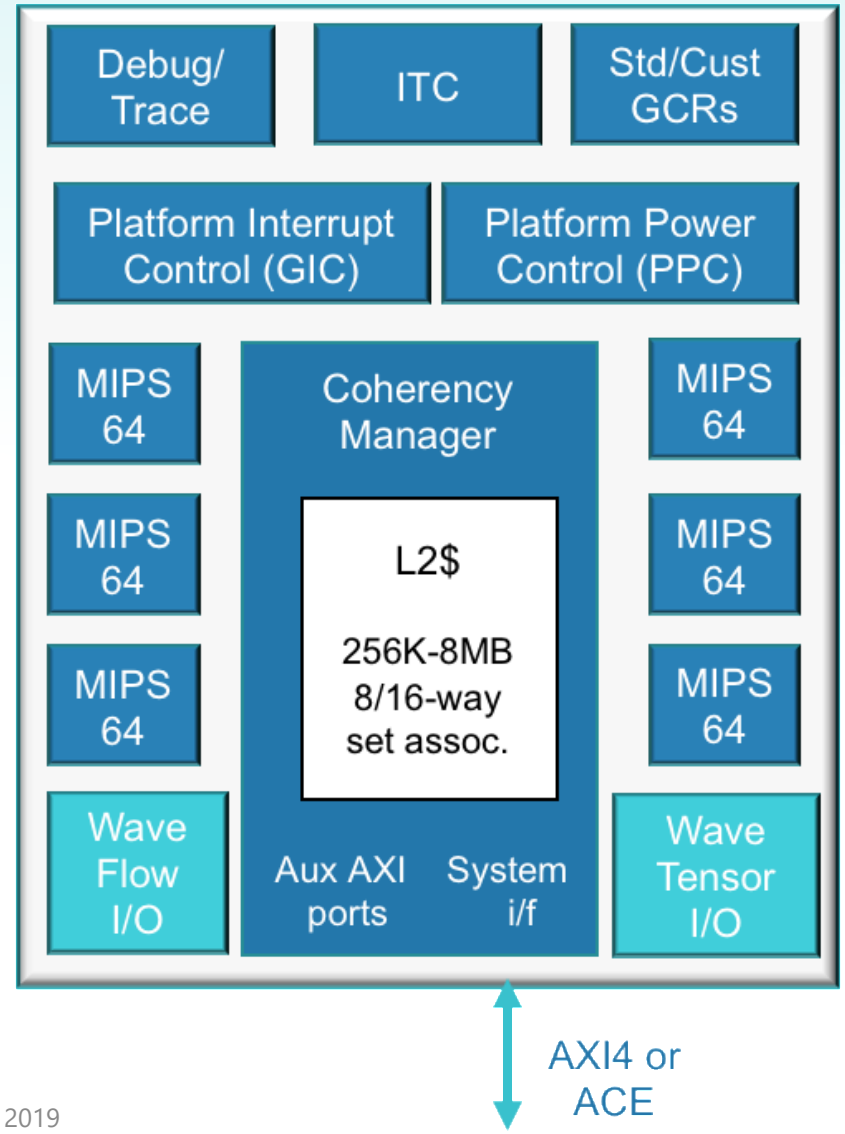

Wave's TritonAI™ 64 IP Software Platform

Software Platform:

- Mature IDE & Tools
- Driver Layer for Technology Mapping
- Linux Operating system support/updates
- Abstract AI Framework calls via WaveRT™ API
- Optimized AI Libraries for:
- CPU/SIMD/WaveFlow/WaveTensor
- TensorFlow-Lite Build support/updates
- Extensible to Edge Training with Full TensorFlow Build

Configurable Hardware Platform:

- MIPS64r6 ISA Cluster
  - 1- 6 cores
  - 1-4 threads/core
  - L1 I/D (32KB-64KB)
  - Unified L2 (256K to 8 Mbytes)
- WaveFlow Tile Array
  - 4 – N Tiles
- WaveTensor Slice Array
  - 1 – N Slices

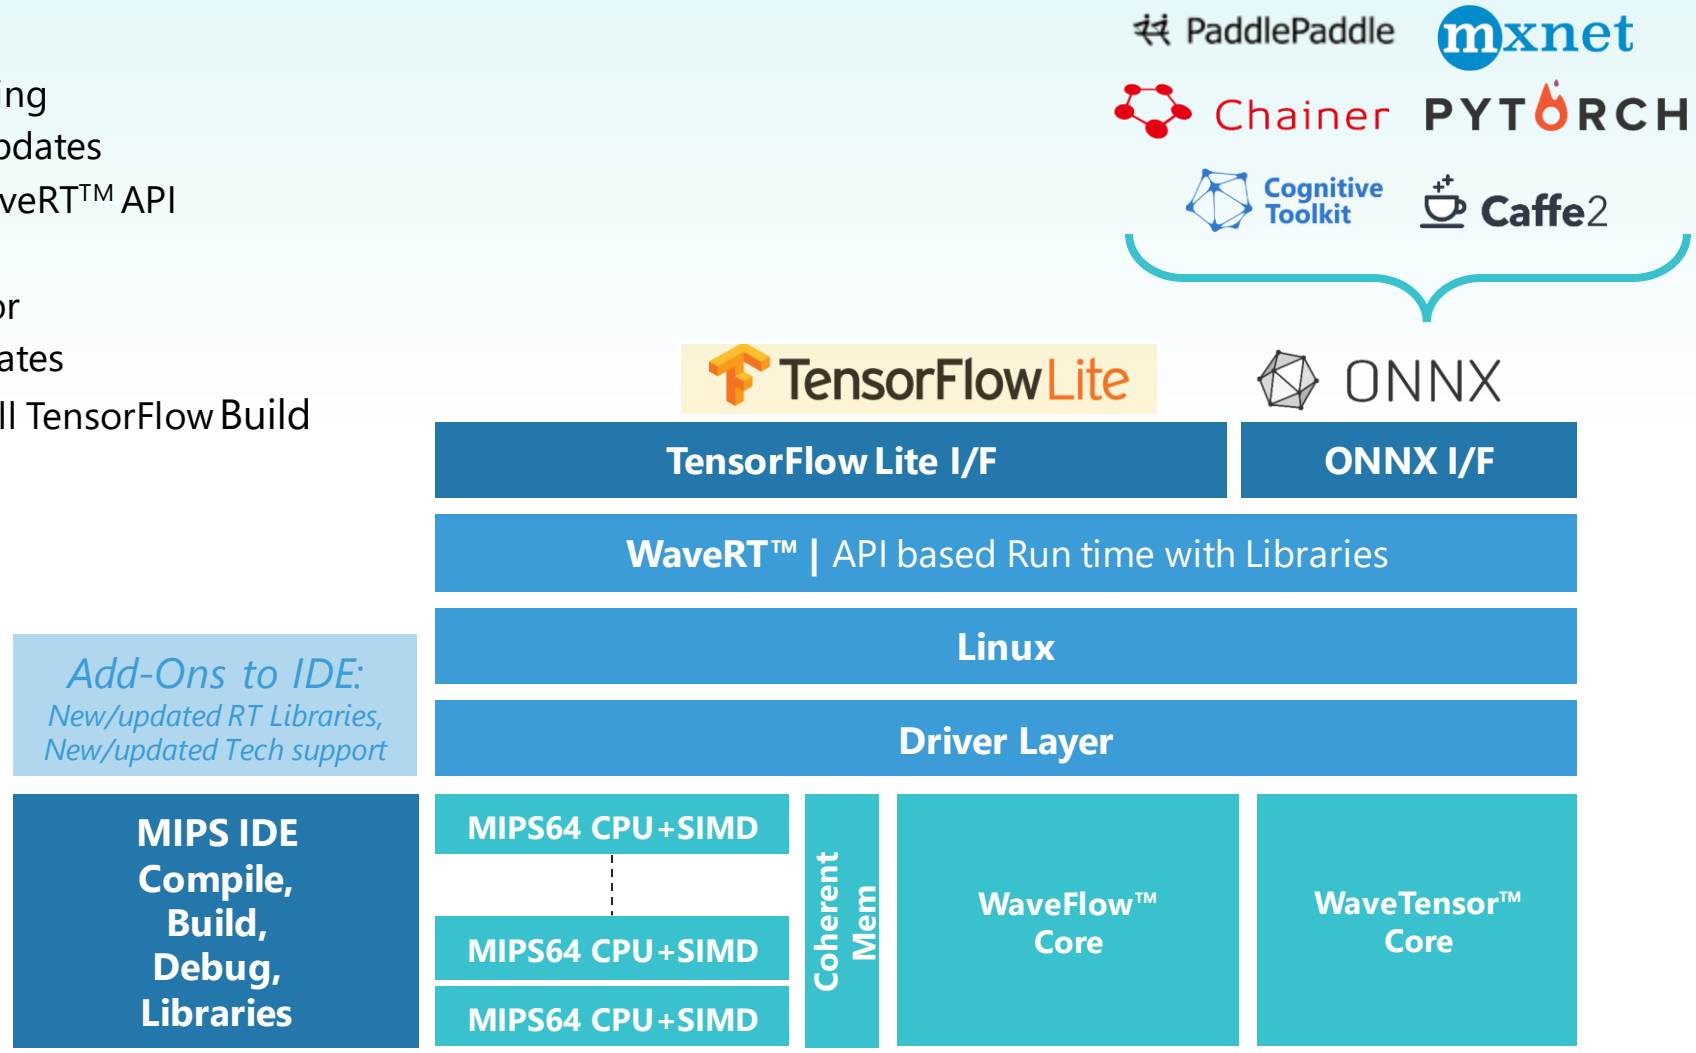


Federated
Learning:
The Next Frontier in
Edge AI

Today's Approach: Centralized Machine Learning

Better ML comes at a cost of collecting data
Most training done in the cloud. i.e. Send your data to the cloud.

Diminished Privacy

- Where is your data?
- Who has access to your data?

Incompatible with Banks, Insurance, Military, Health sectors

Latency Problems

- Most access technologies are asymmetric

High Communications Costs

Federated learning uses training at the edge to refine the global model

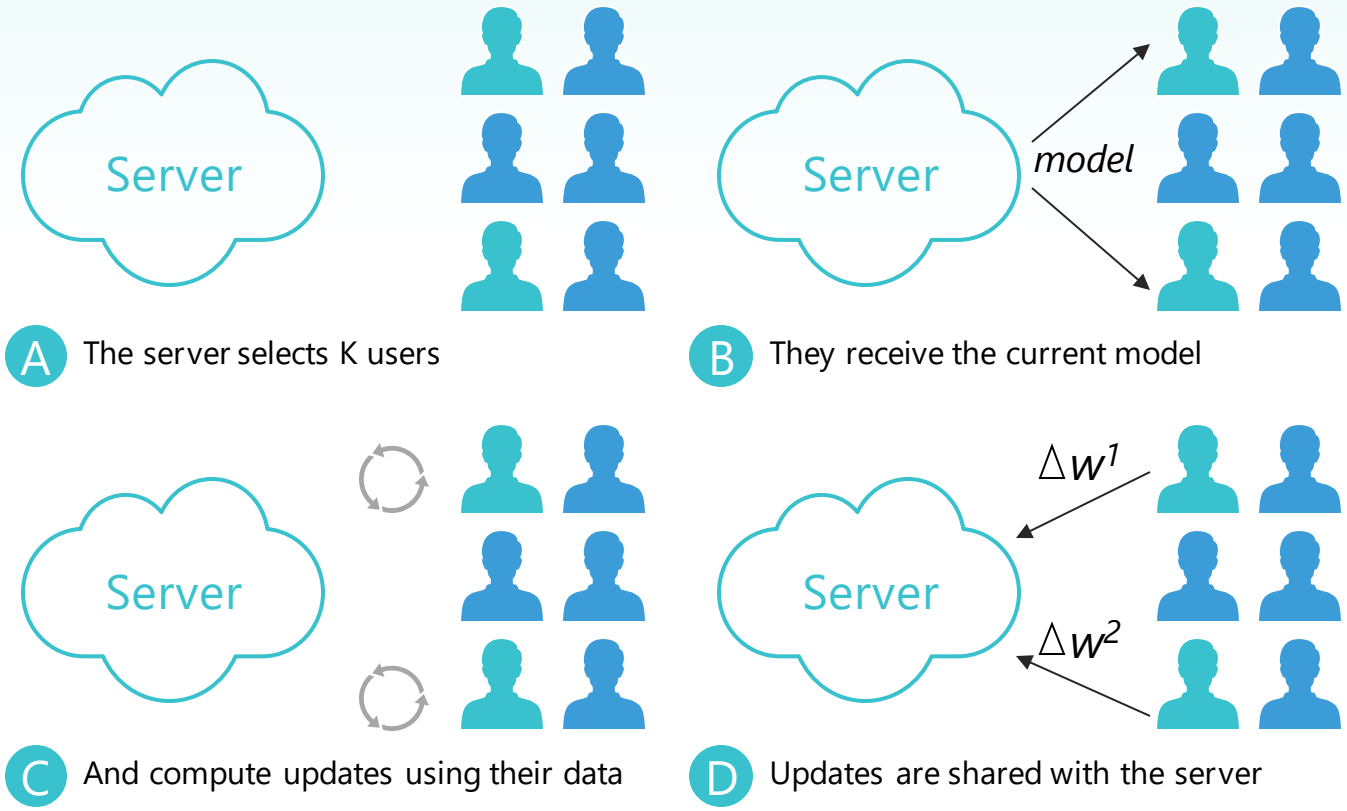

1. Server selects a group of users
2. Users receive copy of central model
3. Users update model based on local data ("Training at the Edge")
4. Updates are shared with the server

   (User data remains private)
5. Server aggregates the changes and updates the central model

Source: Florian Hartmann, Google AI

Federated Learning: Edge AI with Data Privacy

Benefits & Use Cases:

- Transfer learning using local data at edge
- Edge data remains private
- Social networking applications
- Intelligent transportation systems that help increase passenger & pedestrian safety + traffic flow

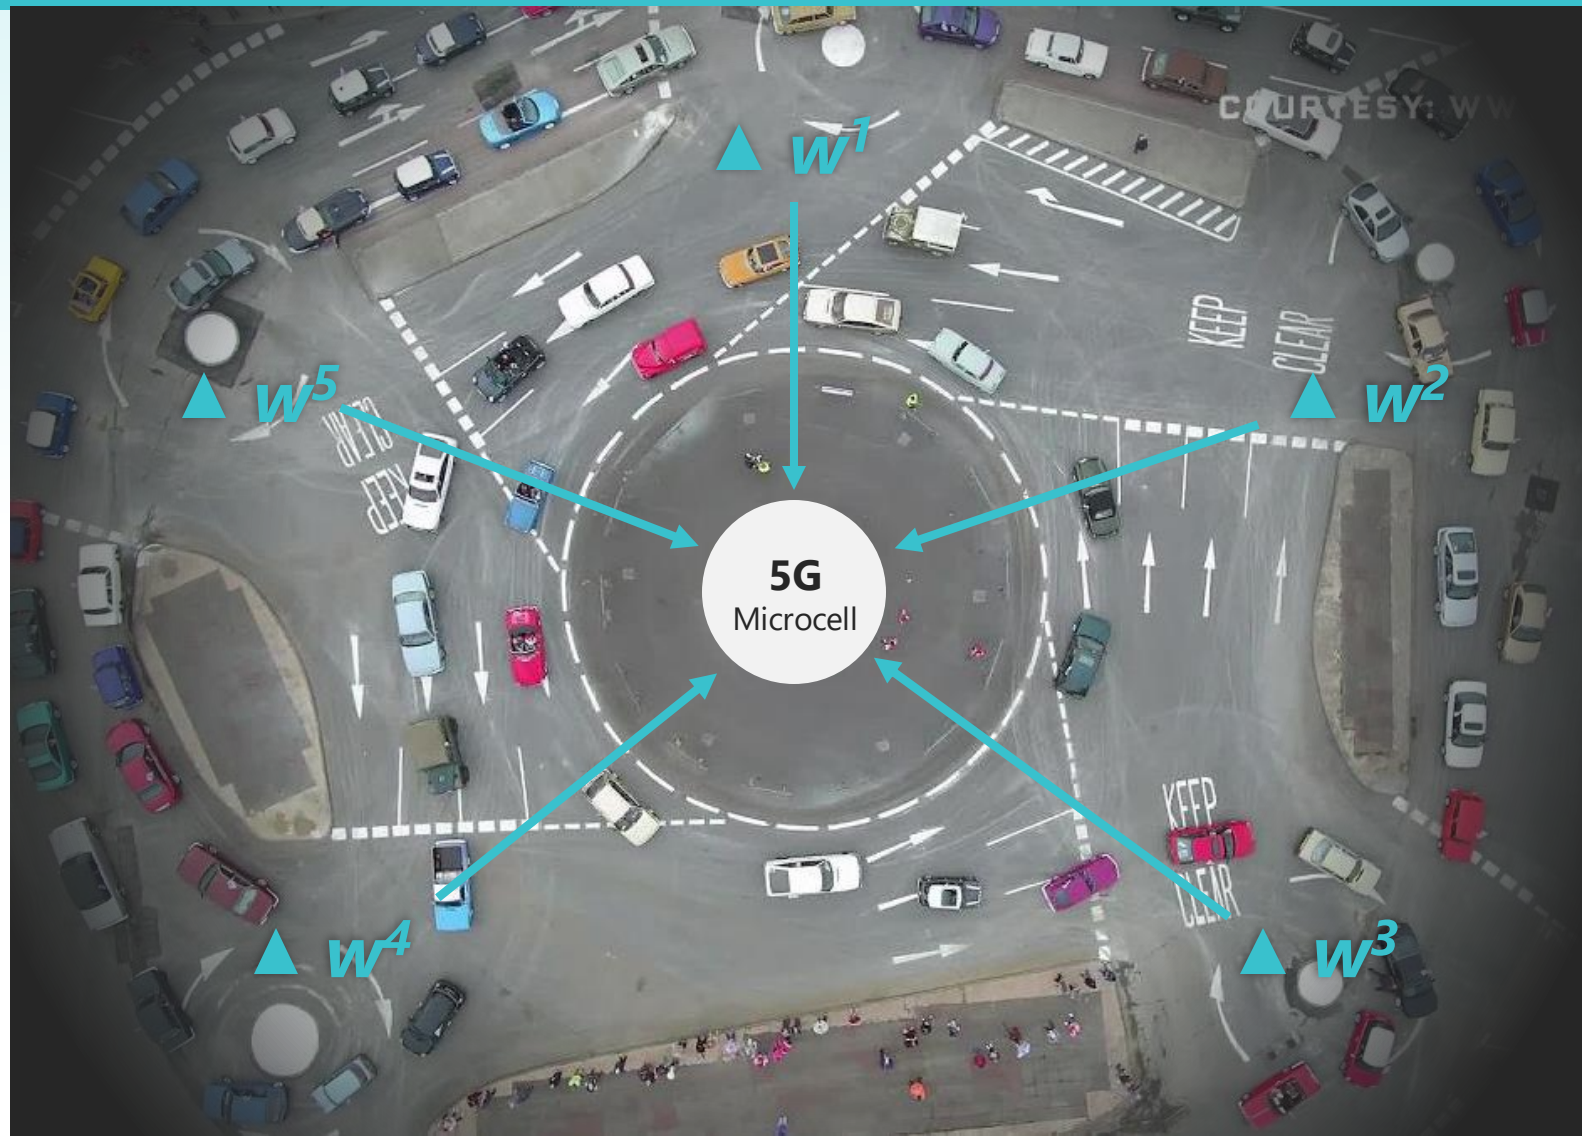

Federated Learning: Edge AI with Data Privacy

Federated learning uses training at the edge
to refine a global master model

Benefits & Use Cases:

- Transfer learning using local data at edge
- Edge data remains private
- Social networking applications
- Intelligent transportation systems that help increase passenger & pedestrian safety + traffic flow

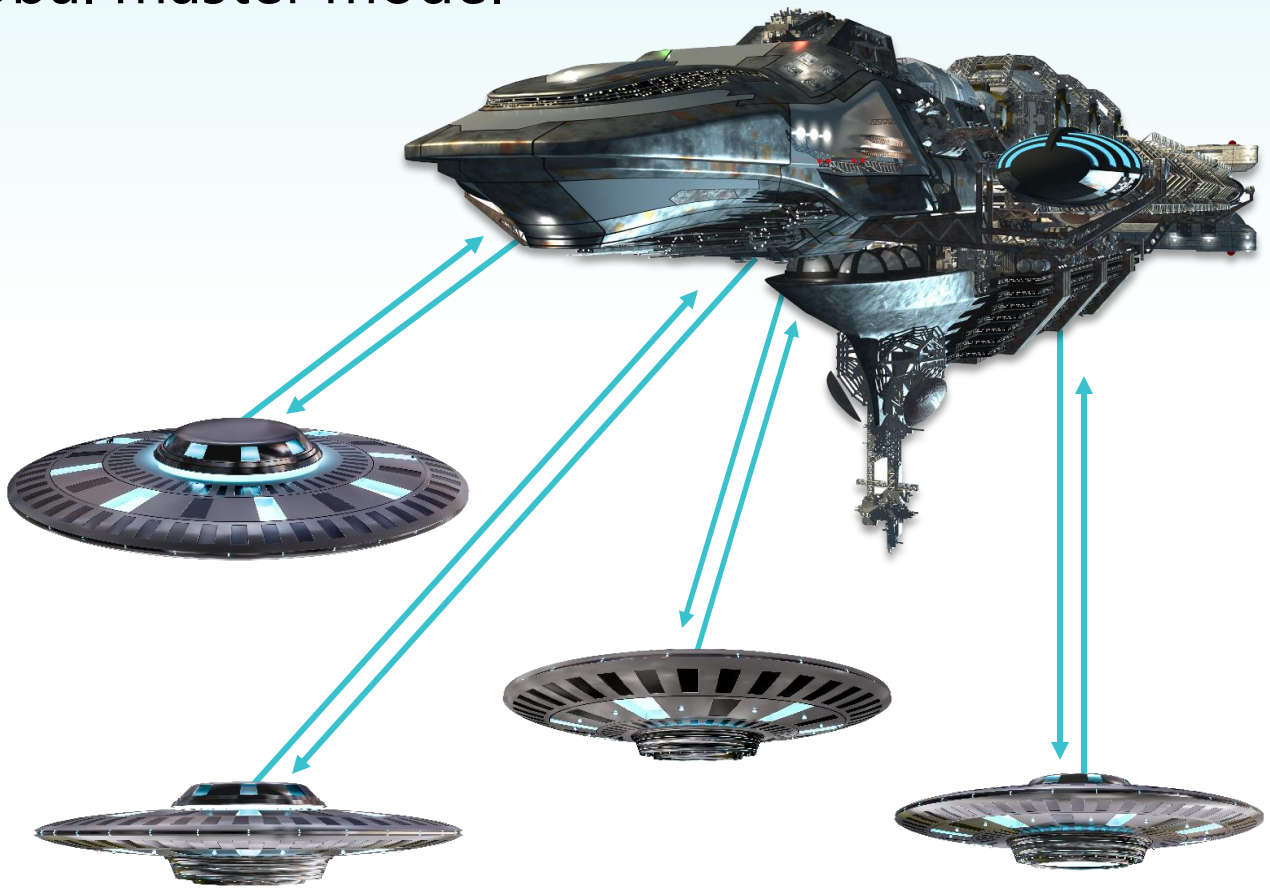

Wave is Driving Training at the Edge with bfloat16

Training into Two Camps:
Float16 or bfloat16 datatypes

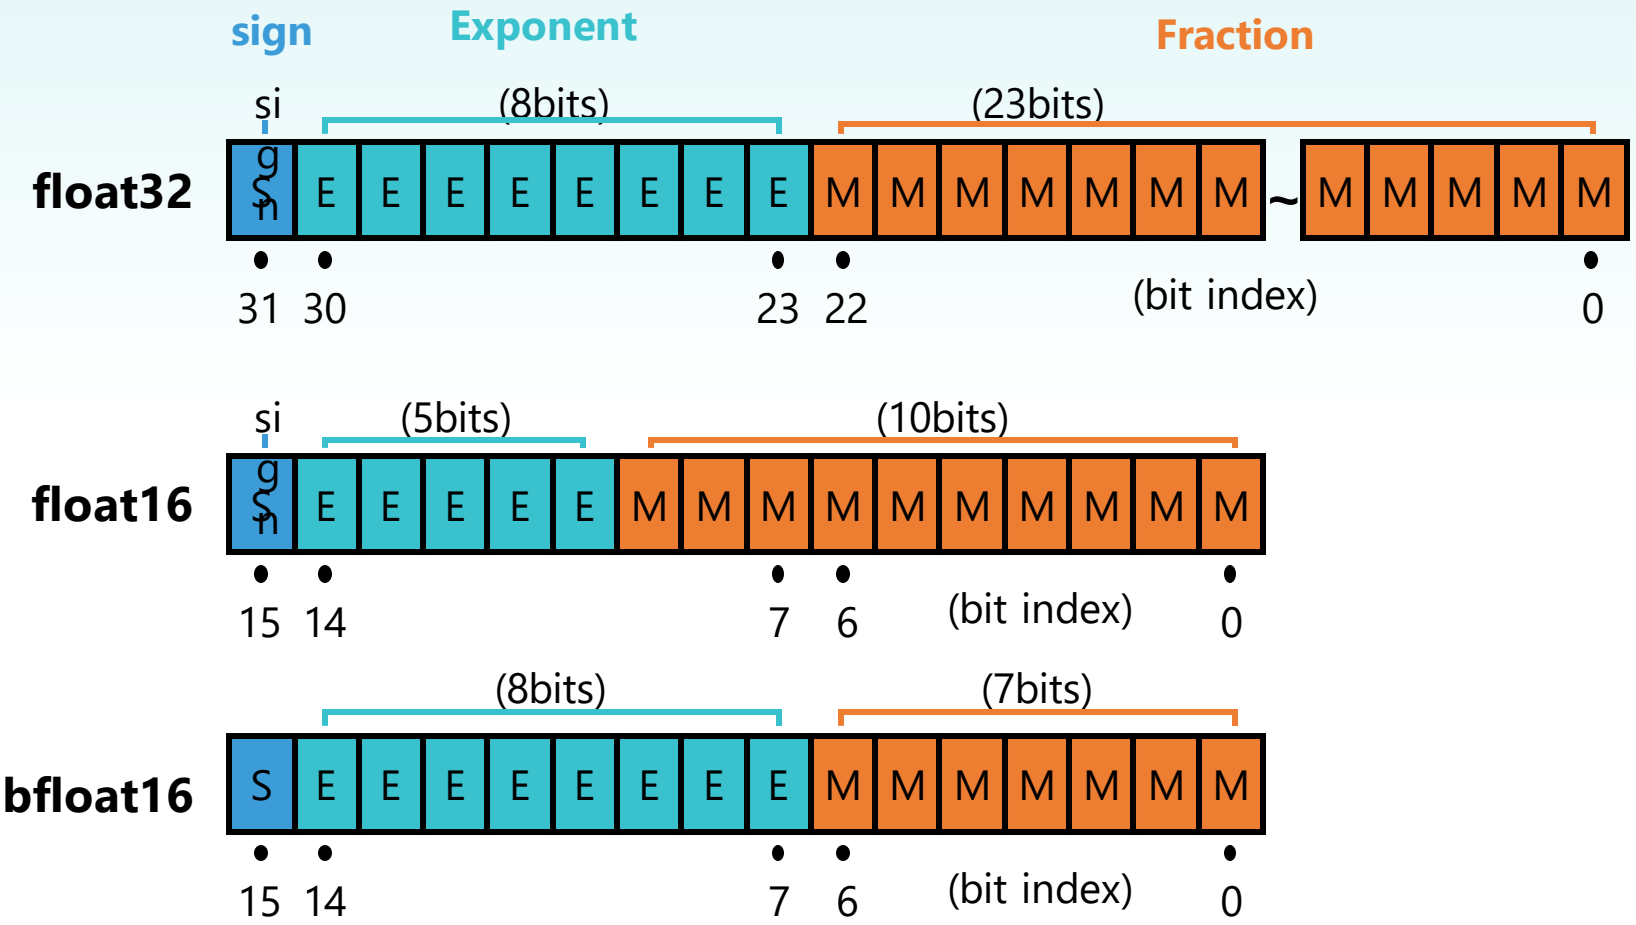

Edge Training Development

- Training Stacks for
  - Federated Learning at the Edge
  - Transfer Learning at the edge
  - Local or personalized models
- Full TensorFlow Build
  - WaveRT API Ext for Training
  - Optimized SIMD FP32 & bfloat16 eigen libraries
  - Deploy training at the edge

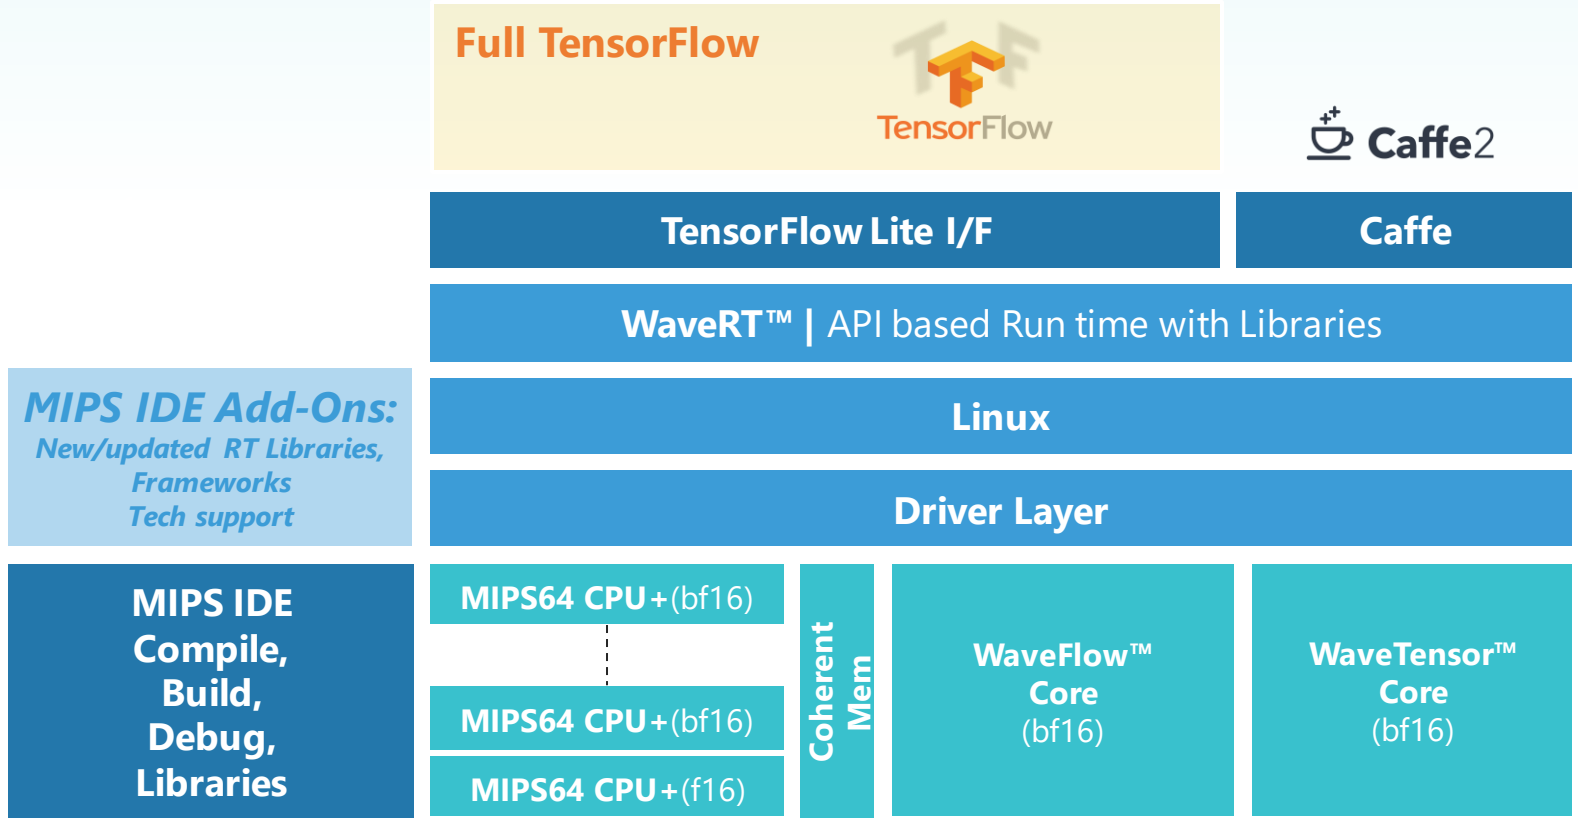

Summary

Wave's TritonAI™ Platform Drives Inferencing to the Edge

Wave's TritonAI™ Platform is a configurable, scalable & programmable offering customers' efficiency, flexibility and AI investment protection

Wave will enable "Training at the edge" with next-gen MIPS AI processor bfloat16 architectures

WAVE COMPUTING

Thank You

If you have questions or would like more information,
visit www.wavecomp.ai


@wavecomputing


https://www.linkedin.com/company/wave-computing


https://www.facebook.com/WaveComp/

MIPS-Classic Cores

Presented by Yuri Panchul
MIPS Open Technical Lead
MIPS Open Meetup in Moscow
April 15, 2019

MIPS IP core product portfolio

1. Three classes of CPU Cores
2. I-Class brings the ultimate differentiation
3. Next 12-months focus on I-Class

P-Class

High single thread performance via OoO multi-issue design.
High data throughput apps

32-bit

P5600

MIPS32, multi-issue, 1-6 cores
VZ, 128b SIMD/FPU, 40b PA

P6600

MIPS64, multi-issue, 1-6 cores
VZ, 128b SIMD/FPU

I-Class

Performance/cost balanced
Efficiency via multi-threading
AP & RT embedded apps

64-bit

I6400

MIPS64, dual issue, 1-6 cores
SMT, VZ, SIMD/FPU

I6500

Many core/multi-cluster
Hetero cluster/multi-cluster
I6400 + DSPRAM, ITC, LL ports

I6500-F

FuSa version I6500
Safety Element out of Context (SEooC) design

32-bit

interAptiv

MIPS32, single issue, 1-4 cores
MT, DSP, FPU, MPU or MMU

I7200

nanoMIPS, dual issue, 1-4 cores
VMT, DSP, MPU/MMU, LL ports

M-Class

Size/Cost optimized CPUs
Class leading perf efficiency
IoT & embedded apps

32-bit

microAptiv

MIPS32/microMIPS,
MCU/MPU, SRAM/L1 cache
DSP, FPU (UCF), SRS

M5100/M5150

microAptiv +
VZ, FPU, Anti-Tamper

M6200/M6250

+30% higher frequency than microAptiv, M51xx
ECC, AXI, 40b PA, APB

MIPS IP Cores - Features Summary

	microAptiv	*M51xx*	M62xx	interAptiv	*I7200*	*I6500/-F*	*P5600*	*P6600*
MIPS Primary ISA	MIPS32 r5	MIPS32 r5	MIPS32 r6	MIPS32 r5	nanoMIPS32	MIPS64 r6	MIPS32 r5	MIPS64 r6
Virtual/Phys Addr Bits	32/32	32/32	32/32	32/32	32/32	48/48	32/40	48/40
FPU	✓ (UC version only)	✓	-	MT	-	MT w/SIMD	Hi Perf w/SIMD	Hi Perf w/SIMD
DSP/SIMD extensions	DSP ASE r2	DSP ASE r2	DSP ASE r2	DSP ASE r2	DSP ASE r2	MSA 128-bit	MSA 128-bit	MSA 128-bit
Virtualization	-	✓	-	-	-	✓	✓	✓
Small code size ISA	microMIPS32	microMIPS32	microMIPS32	MIPS16e2 ASE	nanoMIPS32	-	-	-
Multi-threading	-	-	-	2 VPE, 9 TC	3 VPE, 9 TC	4 VPE	-	-
SuperScalar	-	-	-	-	Dual-issue in order	Dual-issue in order	Multi-issue OoO	Multi-issue OoO
Pipeline stages	5	5	6	9	9	9	16	16
Relative Frequency*	0.6x	0.6x	0.75x	1x	0.95x	0.90x	1.10x	1.10x
SPRAMs (I/D/U)	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / ✓	- / ✓ / -	- / - / -	- / - / -
L1 caches	✓	✓	✓	✓	✓	✓	✓	✓
L2 cache	-	-		✓	✓	✓	✓	✓
Coherent Multi-Core	-	-	-	Up to 4 cores	Up to 4 cores	Up to 6 cores, Up to 64 clusters	Up to 6 cores	Up to 6 cores
Native System Bus I/F	AHB-Lite	AHB-Lite	AXI	OCP 2 or AXI	AXI	AXI or ACE	AXI	AXI

* Relative Frequencies are approximate, are provided for rough guidance only, and will vary to some extent in different process nodes

MIPS technology differentiation

MIPS architecture and IP cores offer powerful, unique capabilities

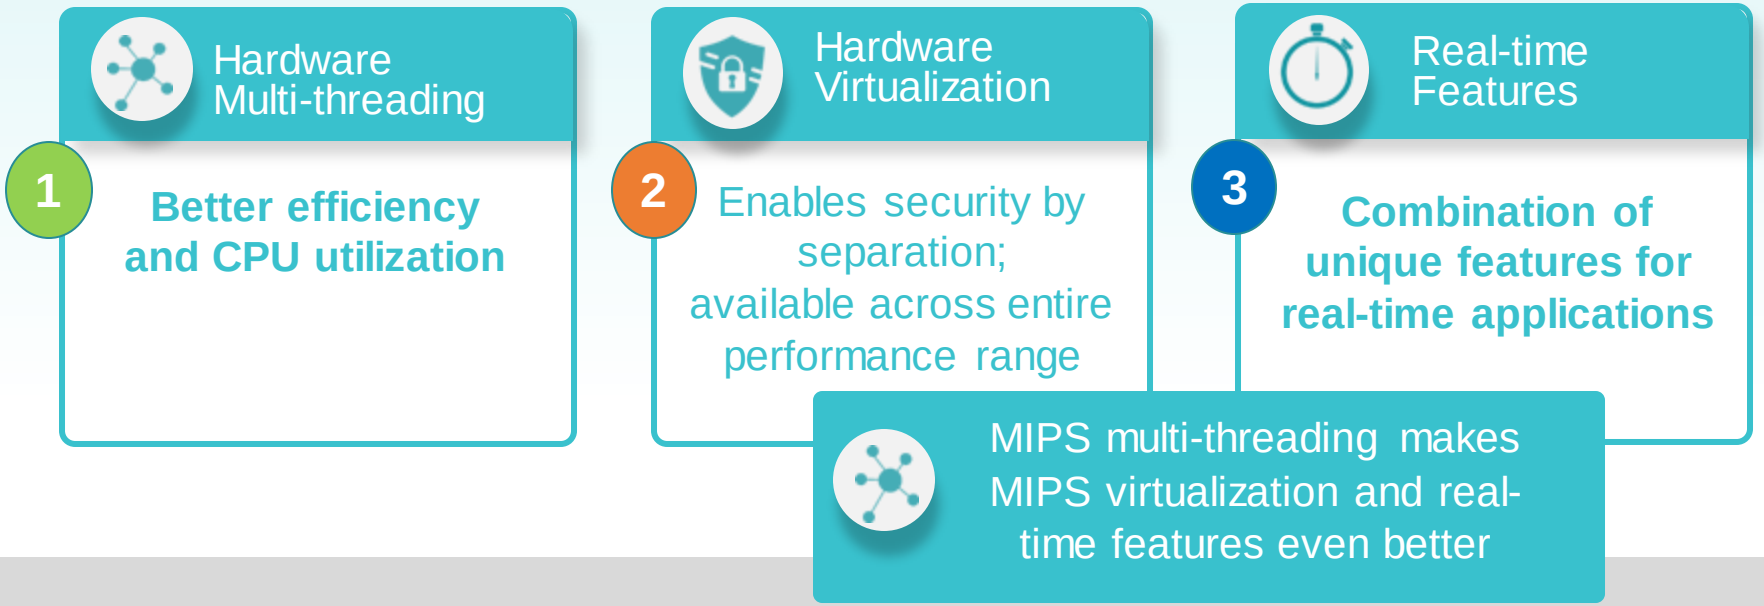

MIPS IP cores

1. Offer leading Power Performance Area (PPA) across the range
2. Provide ultimate scalability: multi-thread, multi-core, multi-cluster
3. Address Functional Safety to ISO 26262 for automotive and IEC 61508 for industrial

WAVE COMPUTING

The I-Class

I6500

I7200

MIPS® I-Class Processor Core Roadmap

MIPS32 and MIPS64 evolution

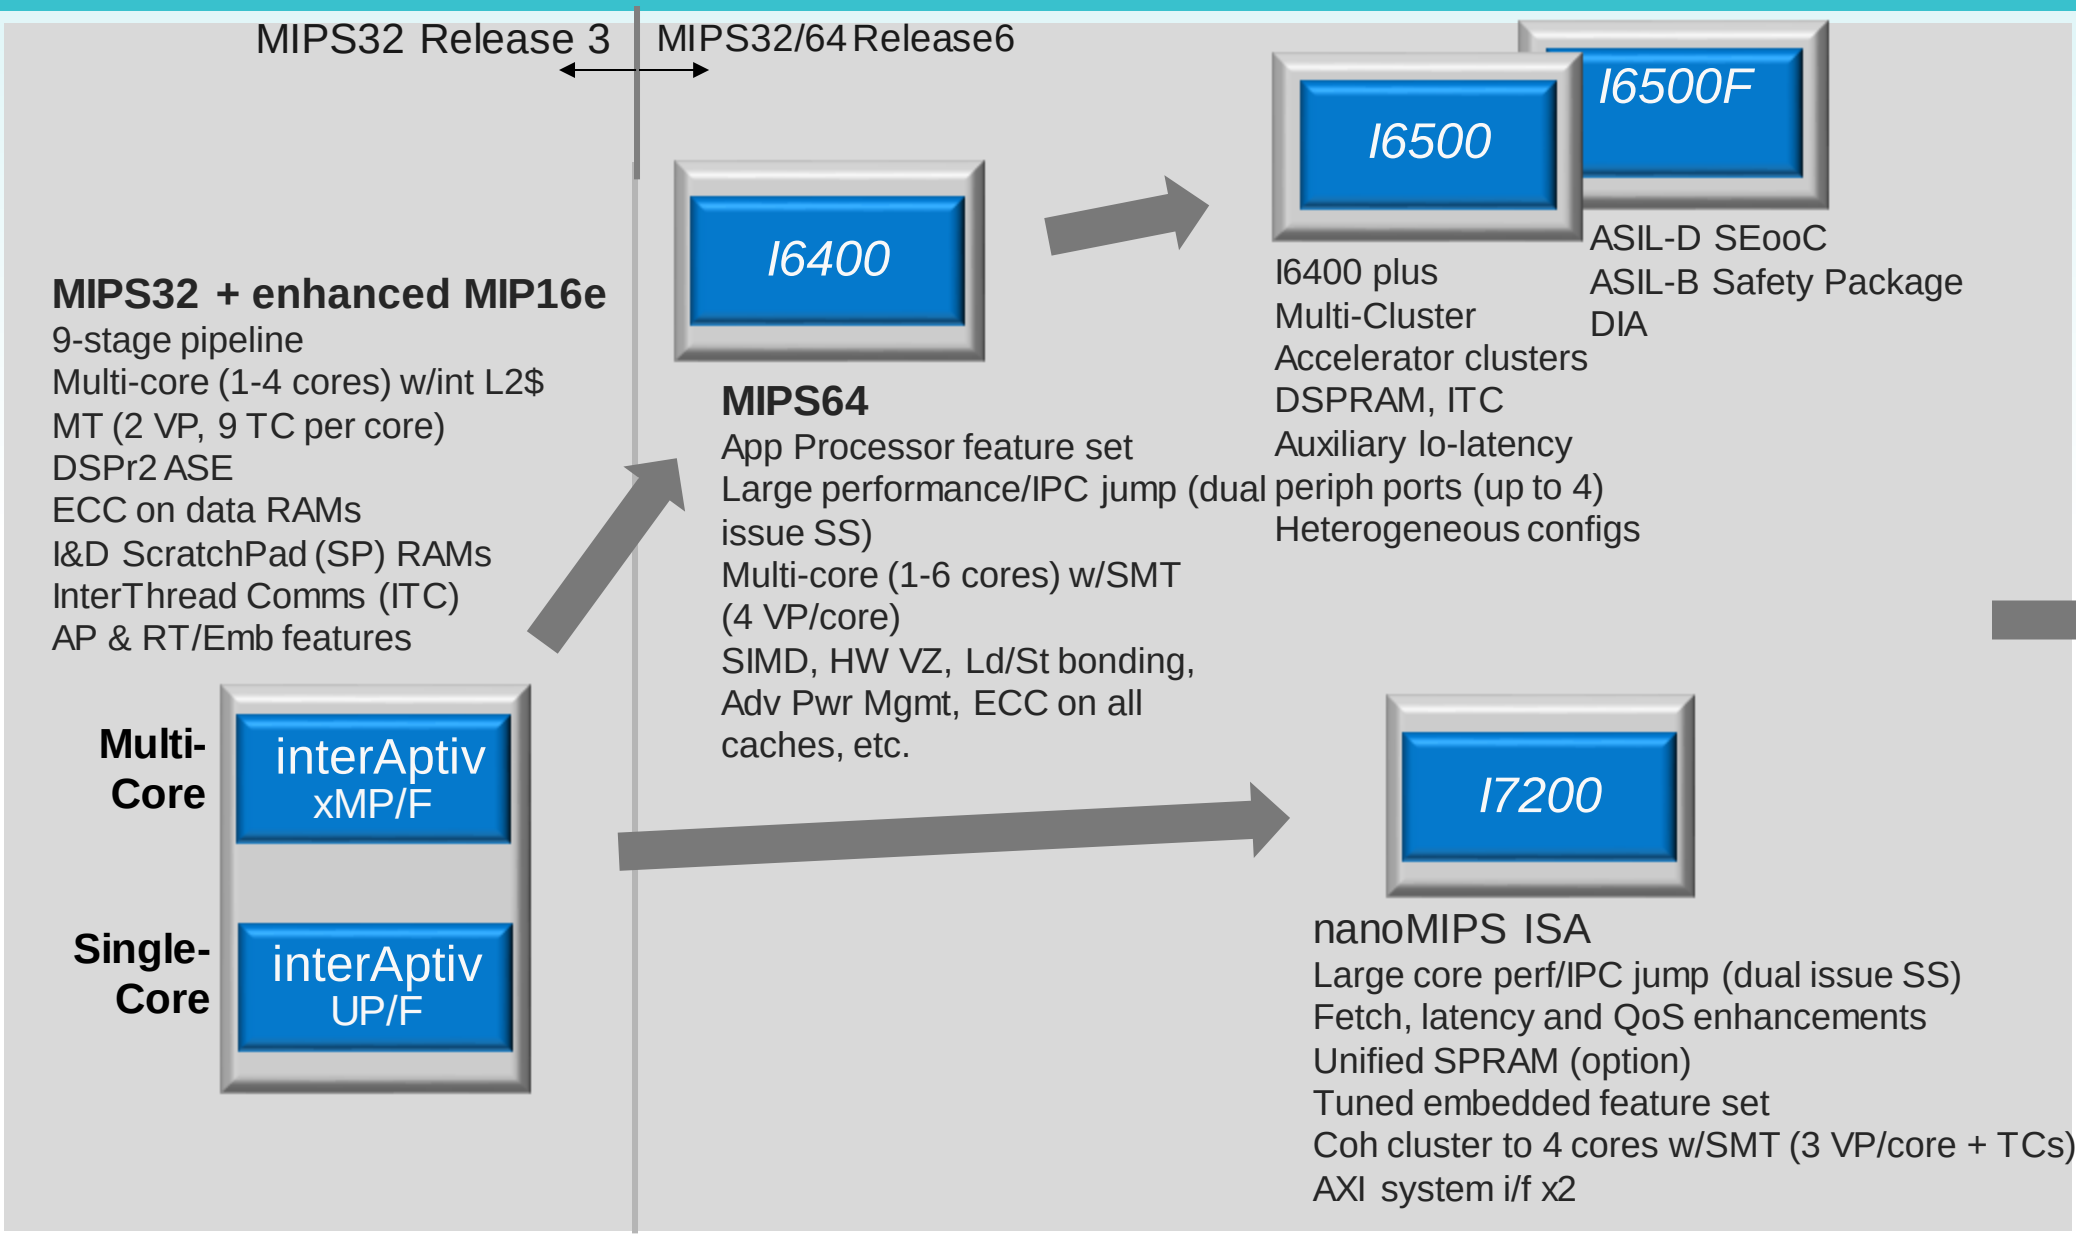

Highly efficient, scalable MIPS Warrior I6400 & I6500

Setting a new standards in mainstream 64-bit processing

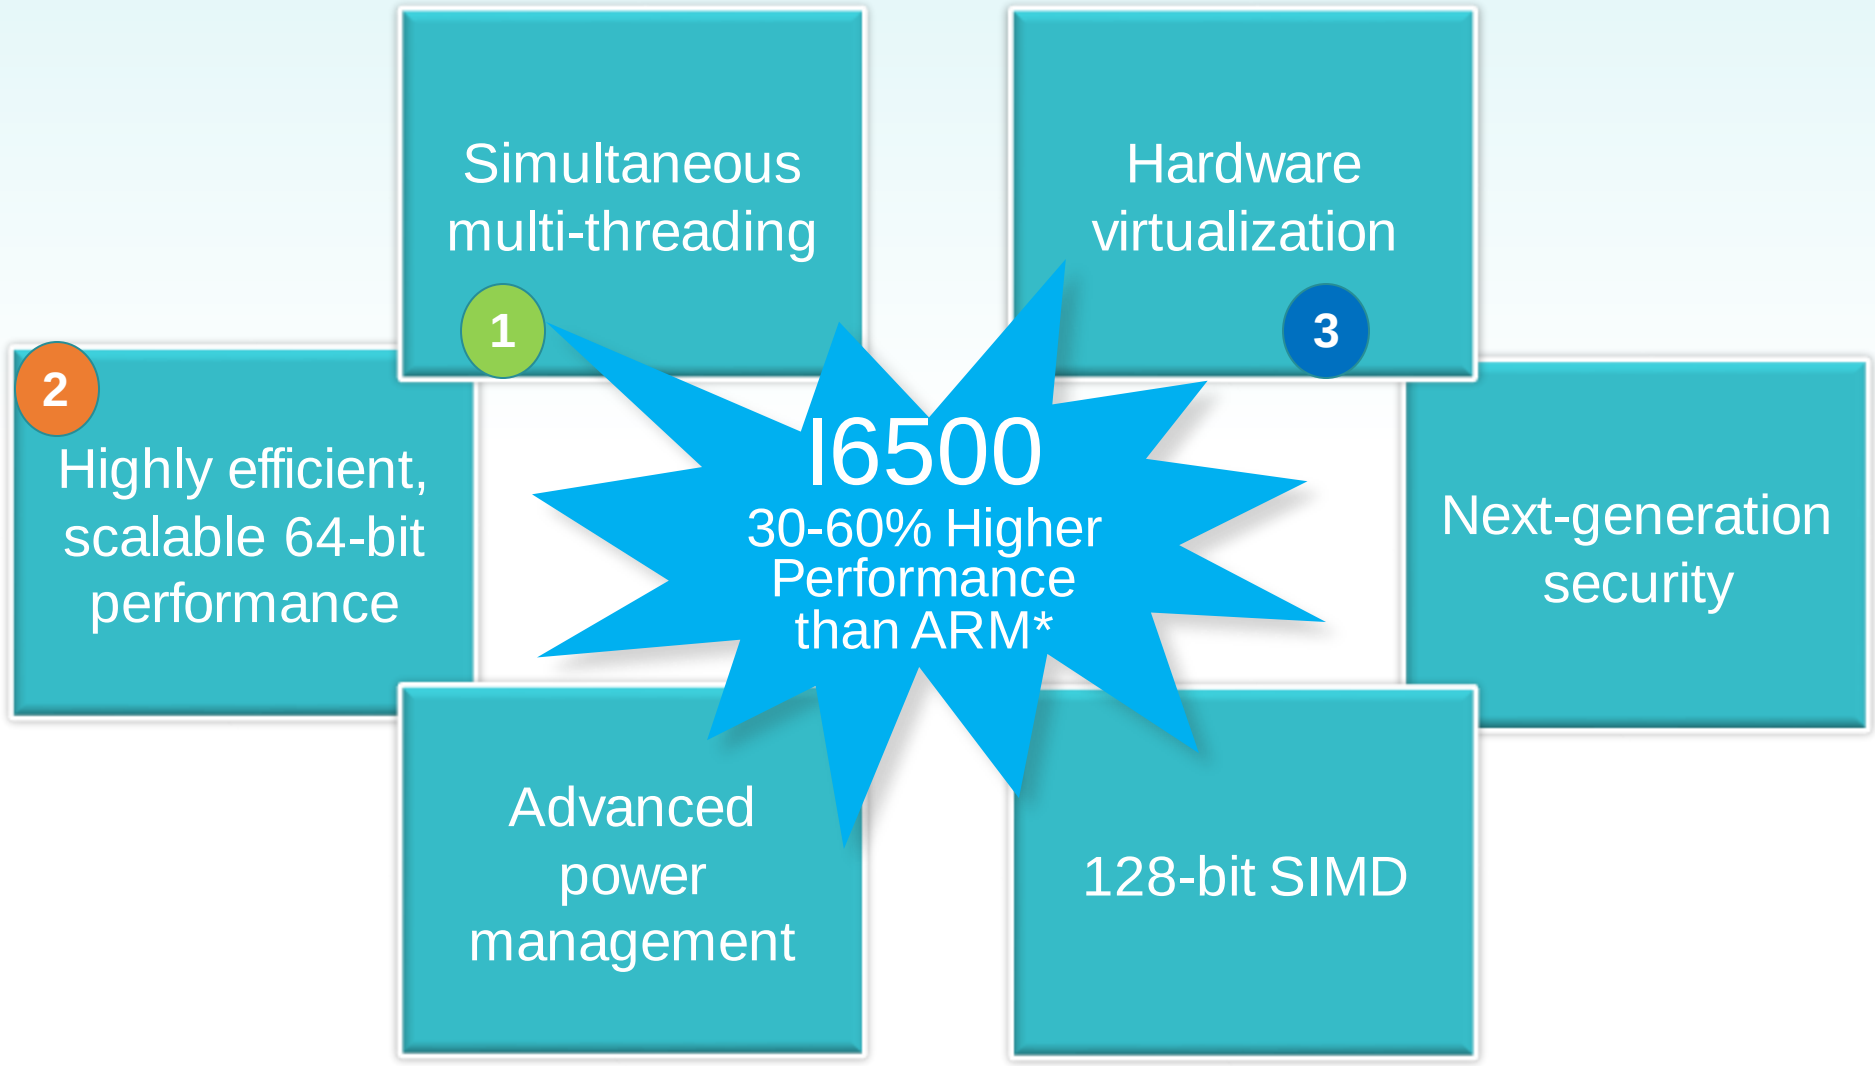

Key Components of a Multi-Core Cluster?

I6500

CPU Cluster : contains multiple CPU Cores, Coherent IO Cores and other supporting CPU Subsystem

MIPS CPU Core: 1 to 4/6 cores / Cluster

Coherency Manager : handles CPU – CPU – IOCU memory transactions coherently

IOCU: 2-8 Coherent IO interface for fast CPU interface to system IO

Cluster Interface: Combination of Industry standard interfaces.
ACE – **A**XI4 **C**oherency **E**xtension

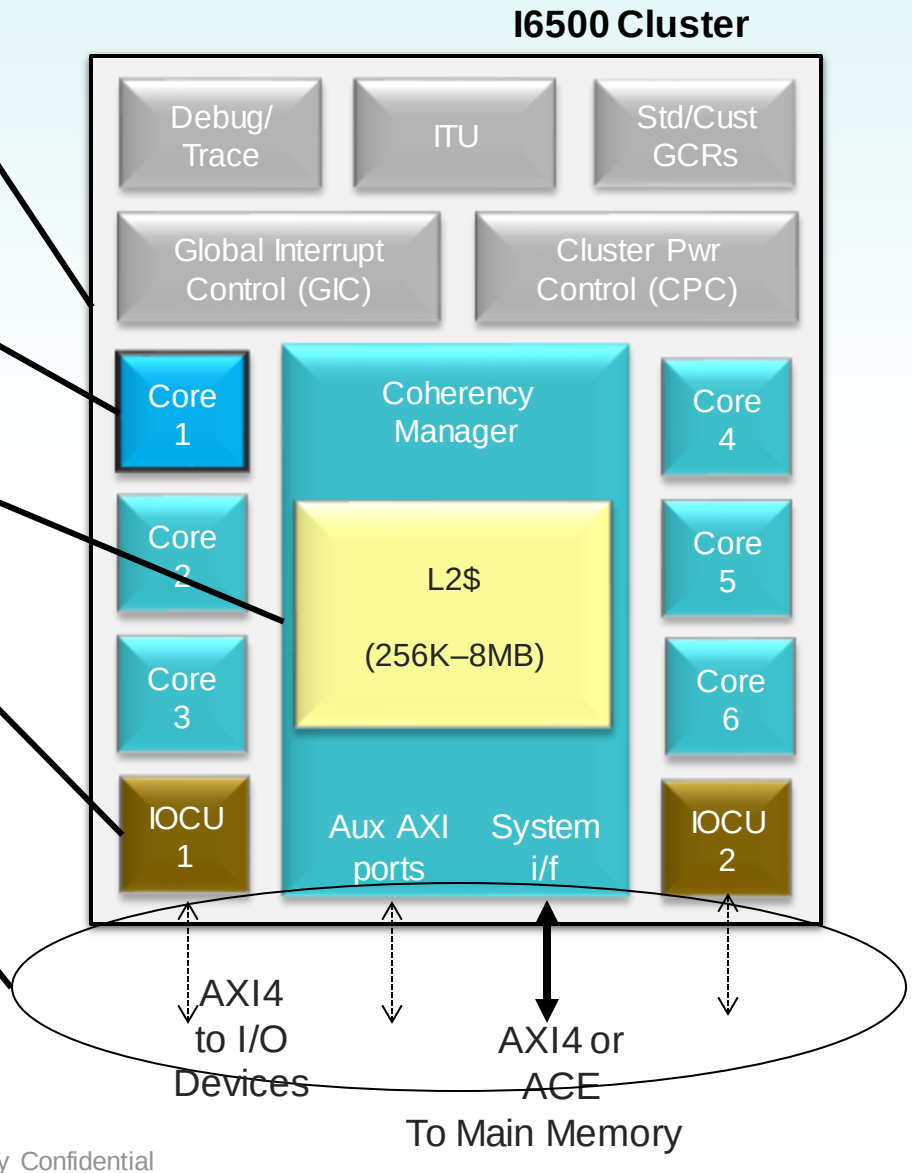

How to build a Multi-Core Cluster?

I6500

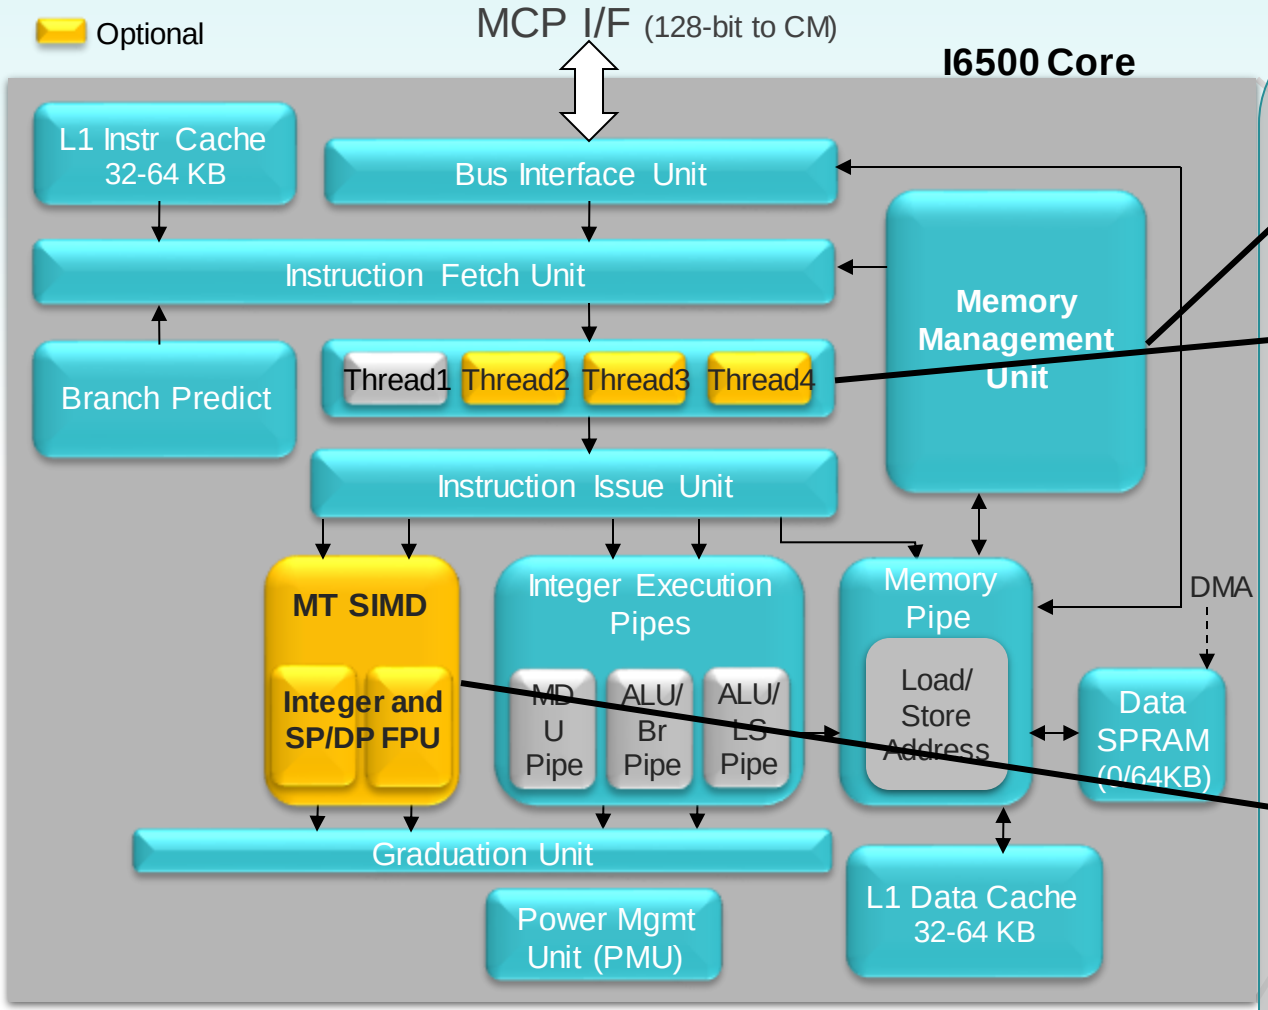

MMU : enable OS like Linux to manage memory efficiently as well as enforcing memory isolation and protection.

SMT (Multi-Threading): up to 4 hardware threads makes I6500 behave like 4-CPUs. Each hardware thread adds ~10% area; 1st thread, boost throughput by as much as 50%

SIMD (Single Instruction Multiple Data): Vector level Arithmetic and Boolean engine working. Ability to execute multiple math operation per cycle.

MIPS I7200 Highlights

- **Multi-threaded multi-core 32-bit processor IP**
- **Designed for high performance embedded systems with real-time requirements**
  - LTE/5G, communications processing
  - Wireless router/gateway processor
  - Data/packet processing

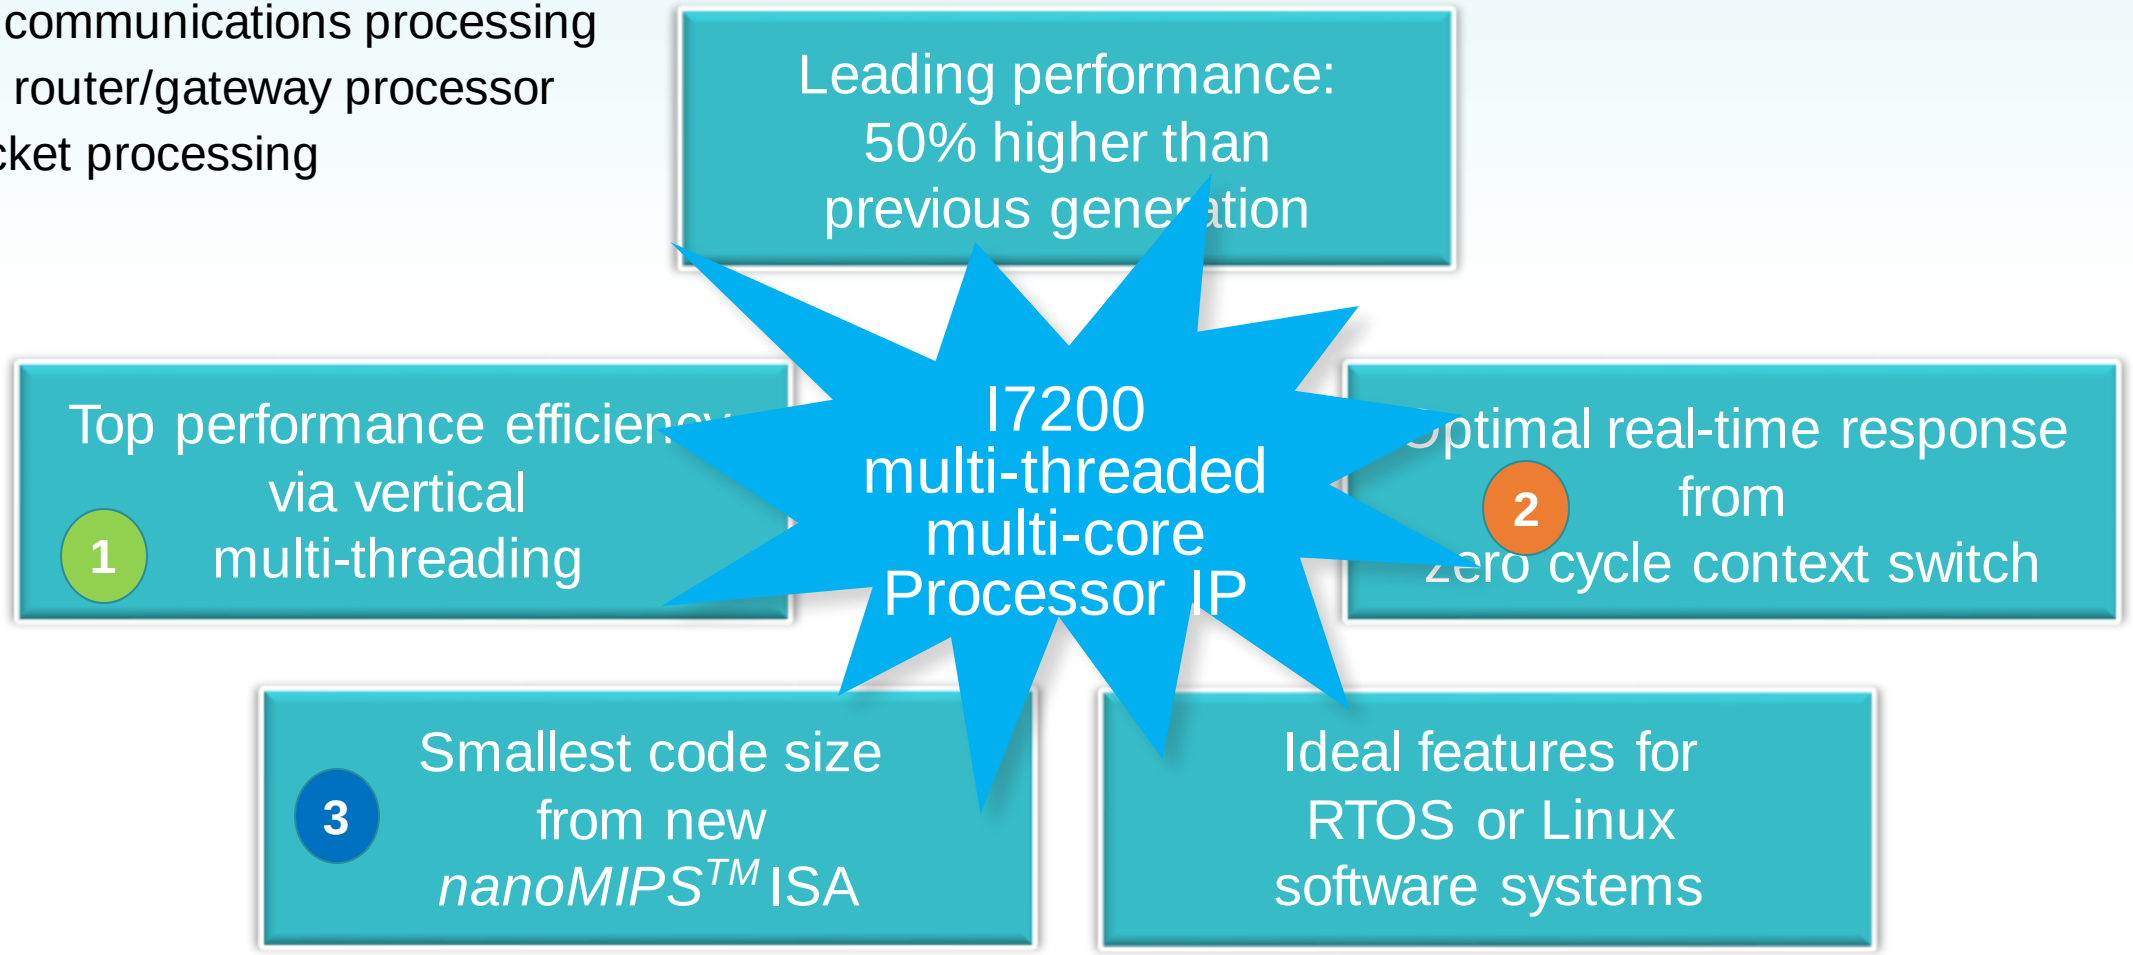

nanoMIPS ISA – advancing small code size

3 Achieves outstanding code density targeting small footprint applications or constraint devices

2 Native ISA for MIPS I7200

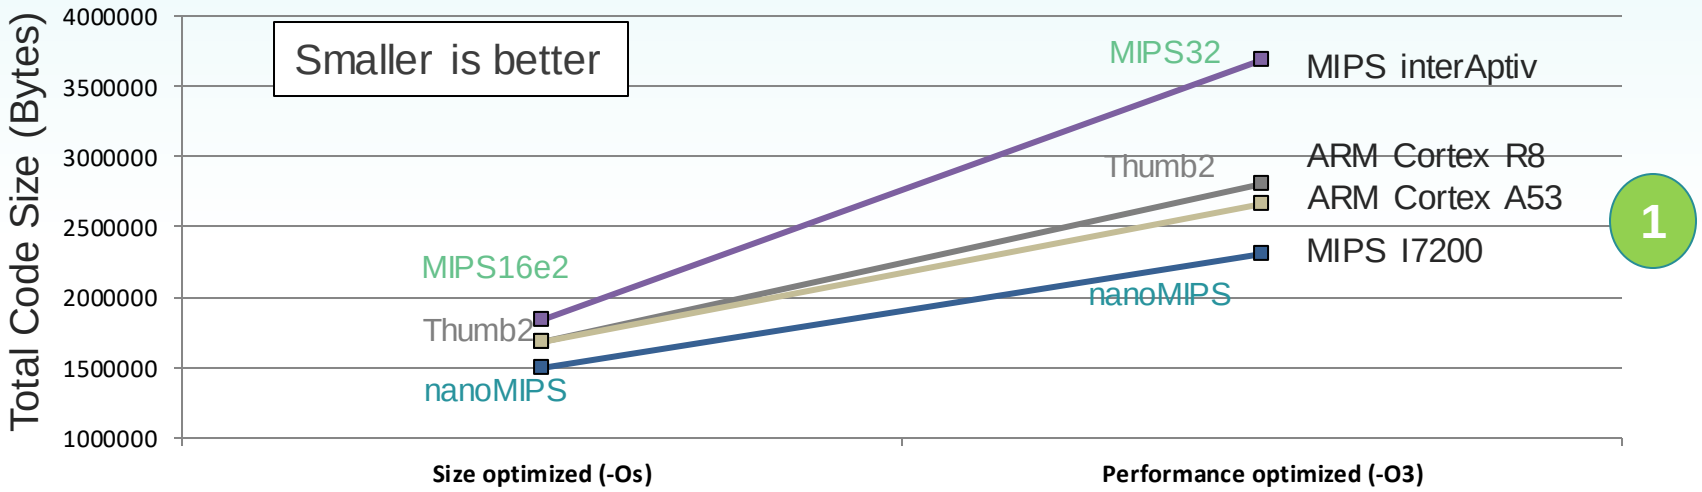

gcc compiler, with indicated optimization target

WAVE COMPUTING

Success Stories

Market specific use-cases

Classical Application Segments

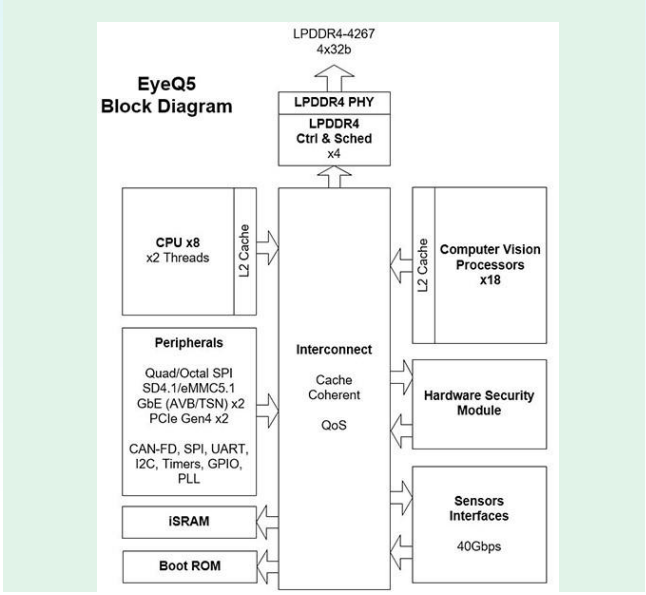

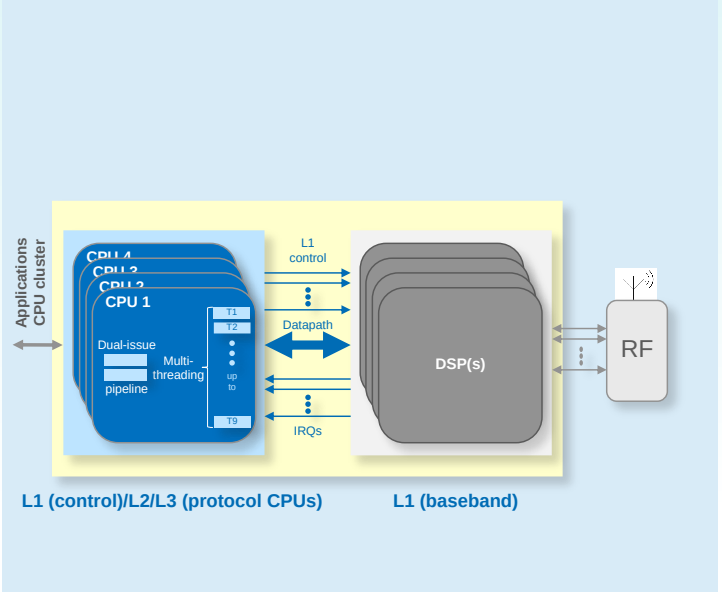

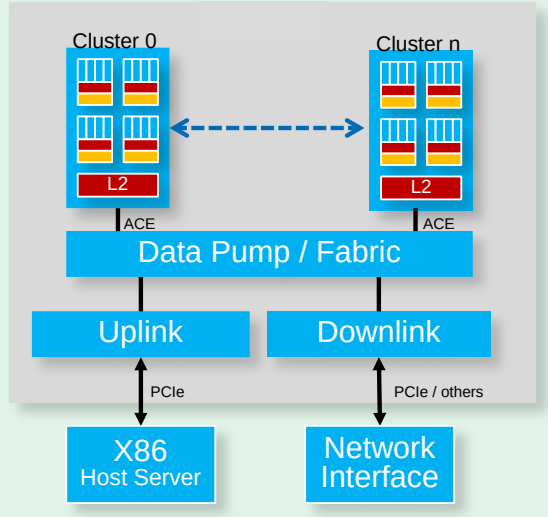

Feature	ADAS	LTE/5G	Datacenter
Application 1	Autonomous driving	Broadband LET modem	Converged Infrastructure. Distributed Workload processing
Value Prop 2	SMT, RT, IPC, VZ	SMT, RT, IPC, low power, small code	SMT, RT, IPC, distributed work load, VZ
Key drivers 3	Heterogeneous, FuSa / SEooC, multi-Cluster	High efficiency class leading 32b uArch High perf small code size ISA Low latency embedded / RT features	Threads / cores / clusters low latency payload processing
Likely Customers	Denso, Renesas, NXP, OnSemi, Samsung, Toshiba, Cypress, STM, Microchip, Bosch,.. Baidu, Argo, …	MTK, UniSoc, Qualcomm, Sony,	Cisco, Huawei, Broadcom, Marvell, Juniper, Google,

WAVE COMPUTING

Application Segment 1
ADAS

I6500-F lead customer: Mobileye ADAS Platform

I6500-F at heart of heterogeneous coherent processing clusters in next-gen EyeQ®5 SoC

Designed to act as central computer for sensor fusion in Fully Autonomous Driving (FAD) vehicles starting 2020

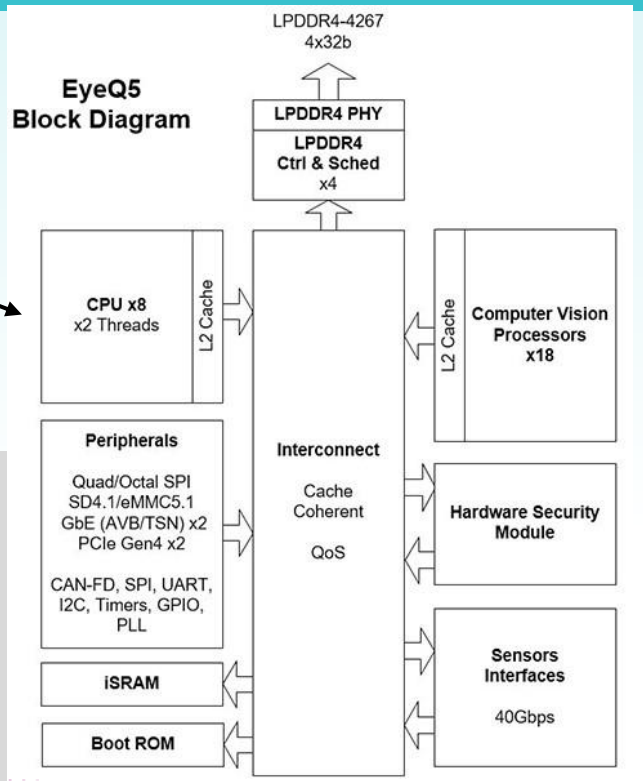

"Our EyeQ®5 SoC will be the most advanced solution of its kind for fully autonomous vehicles which will start rolling out in 2020. The ASIL B(D) features in the I6500-F are key to ensuring our chip achieves the highest level of safety."

2018

– **Elchanan Rushinek, SVP Engineering, Mobileye**

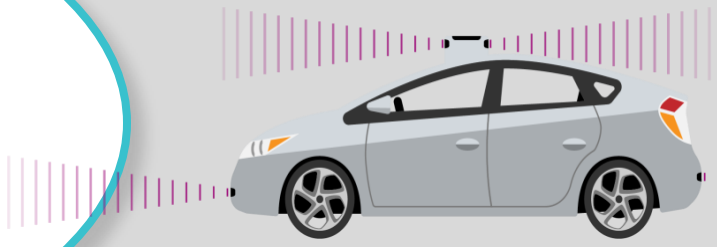

MOBILEYE
An Intel Company

MIPS I-class vs Cortex A7 & A53 performance

Per core, normalized to A53 Values (@ same frequency)

2 Less power or more CPU headroom

1 45+% performance boost for 10% Area

3 I6500 performs best for different workloads

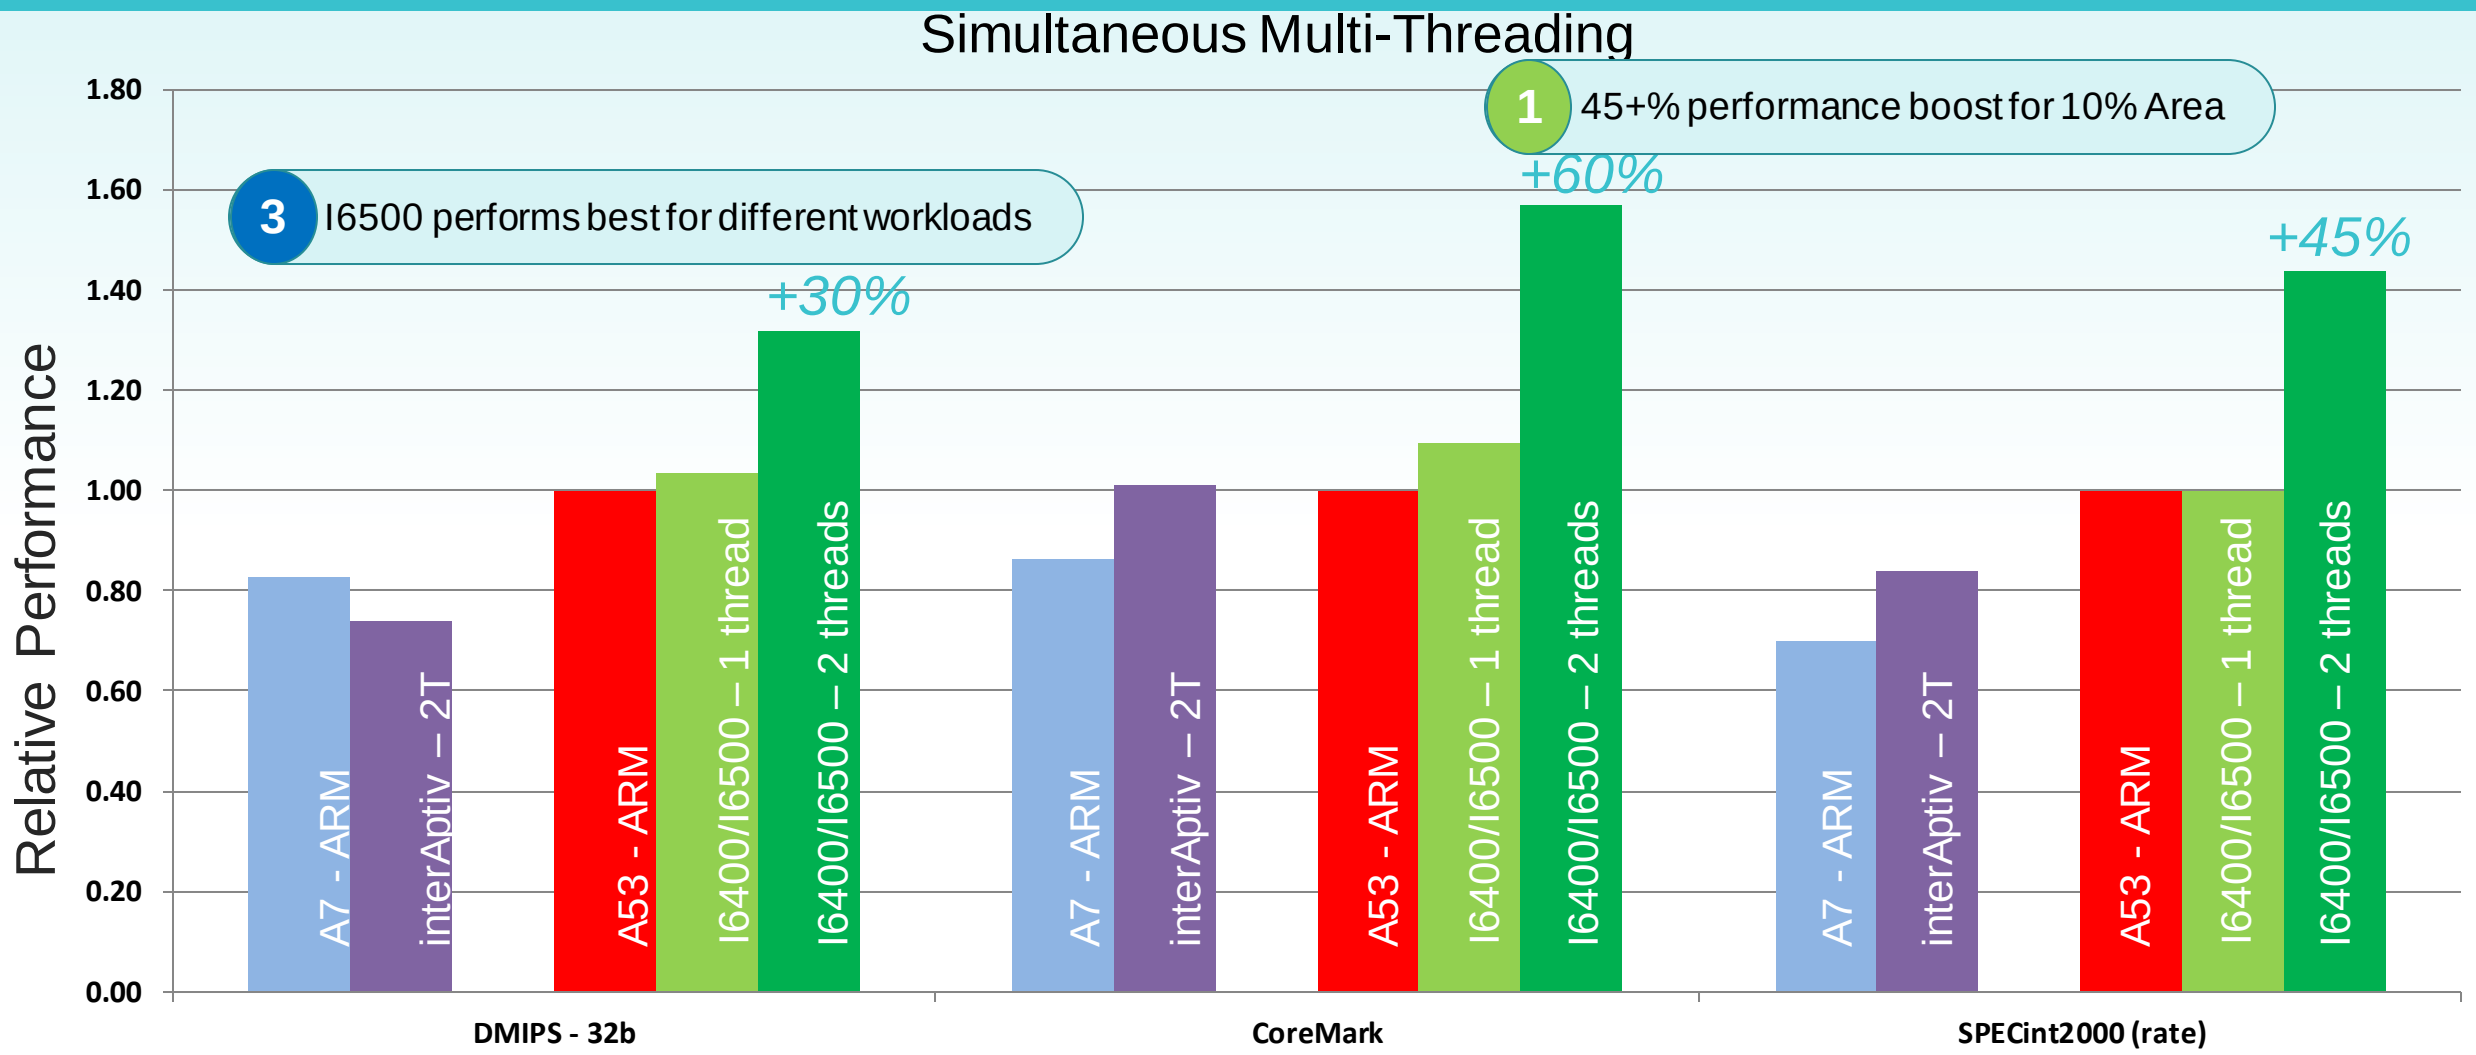

- Based on Cortex A53 data reported by ARM on website and/or in presentation materials, plus benchmarked results on Linaro (HiSilicon Octa A53 Kirin620) with Linux kernel: 3.18.0-linaro-hikey SMP preempt, RFS Debian squeeze, with GCC-based 5.0.0 toolchain. A7 scores are ARM claims. Measured results are lower.
- I6400 results are based on production released RTL, FPGA platform benchmarking and in case of SPEC, 1 enhancement for next release performance models testing

WAVE COMPUTING

Application Segment 2
LTE/5G

Success Story, Multi-threading & Mediatek

I7200

Mediatek 5G LTE modem

LTE

- Mediatek designed MIPS into modem across multiple products
- First version, Helio X30, in production

MIPS Advantage

- **Hardware Multi-Threading (MT)**
  - Fast inter-thread communications
  - Higher performance/high processing efficiency
- **Scalability**
  - 1-4 cores, 2 threads per core
- **Deterministic ; Real-time interrupts**

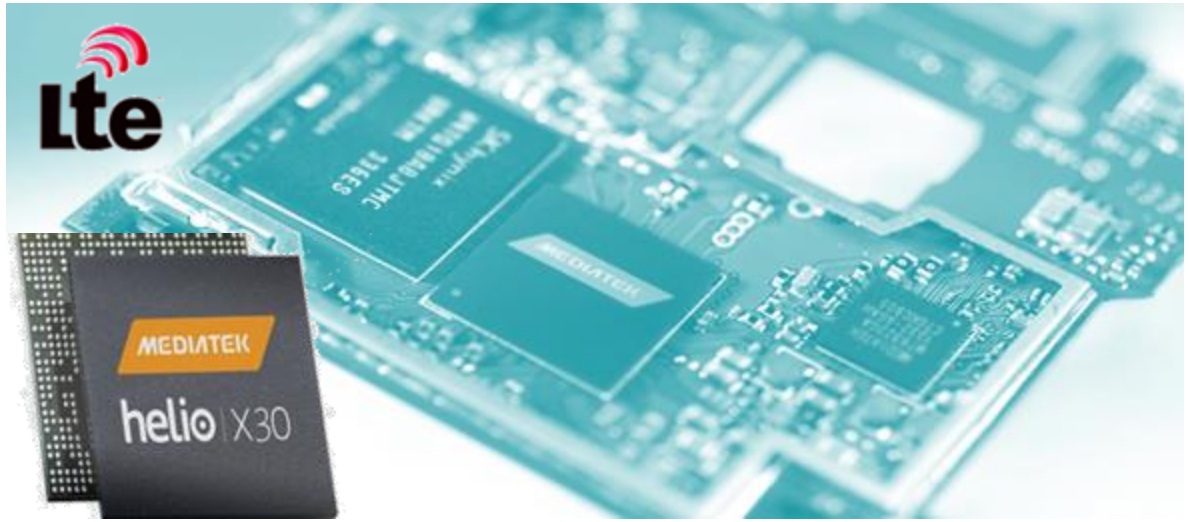

About MIPS in the application

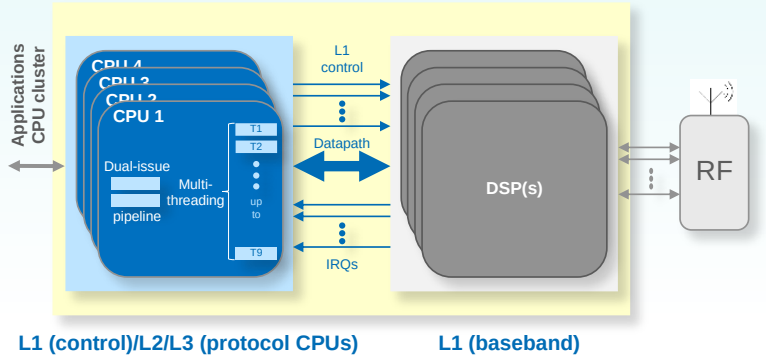

"MIPS CPUs, with their powerful multi-threading capability, offer a combination of efficiency and high throughput for LTE modems that contributes significantly to system performance." ***said Dr. Kevin Jou, SVP and CTO, MediaTek***. **2018**

I7200 Performance Advantage

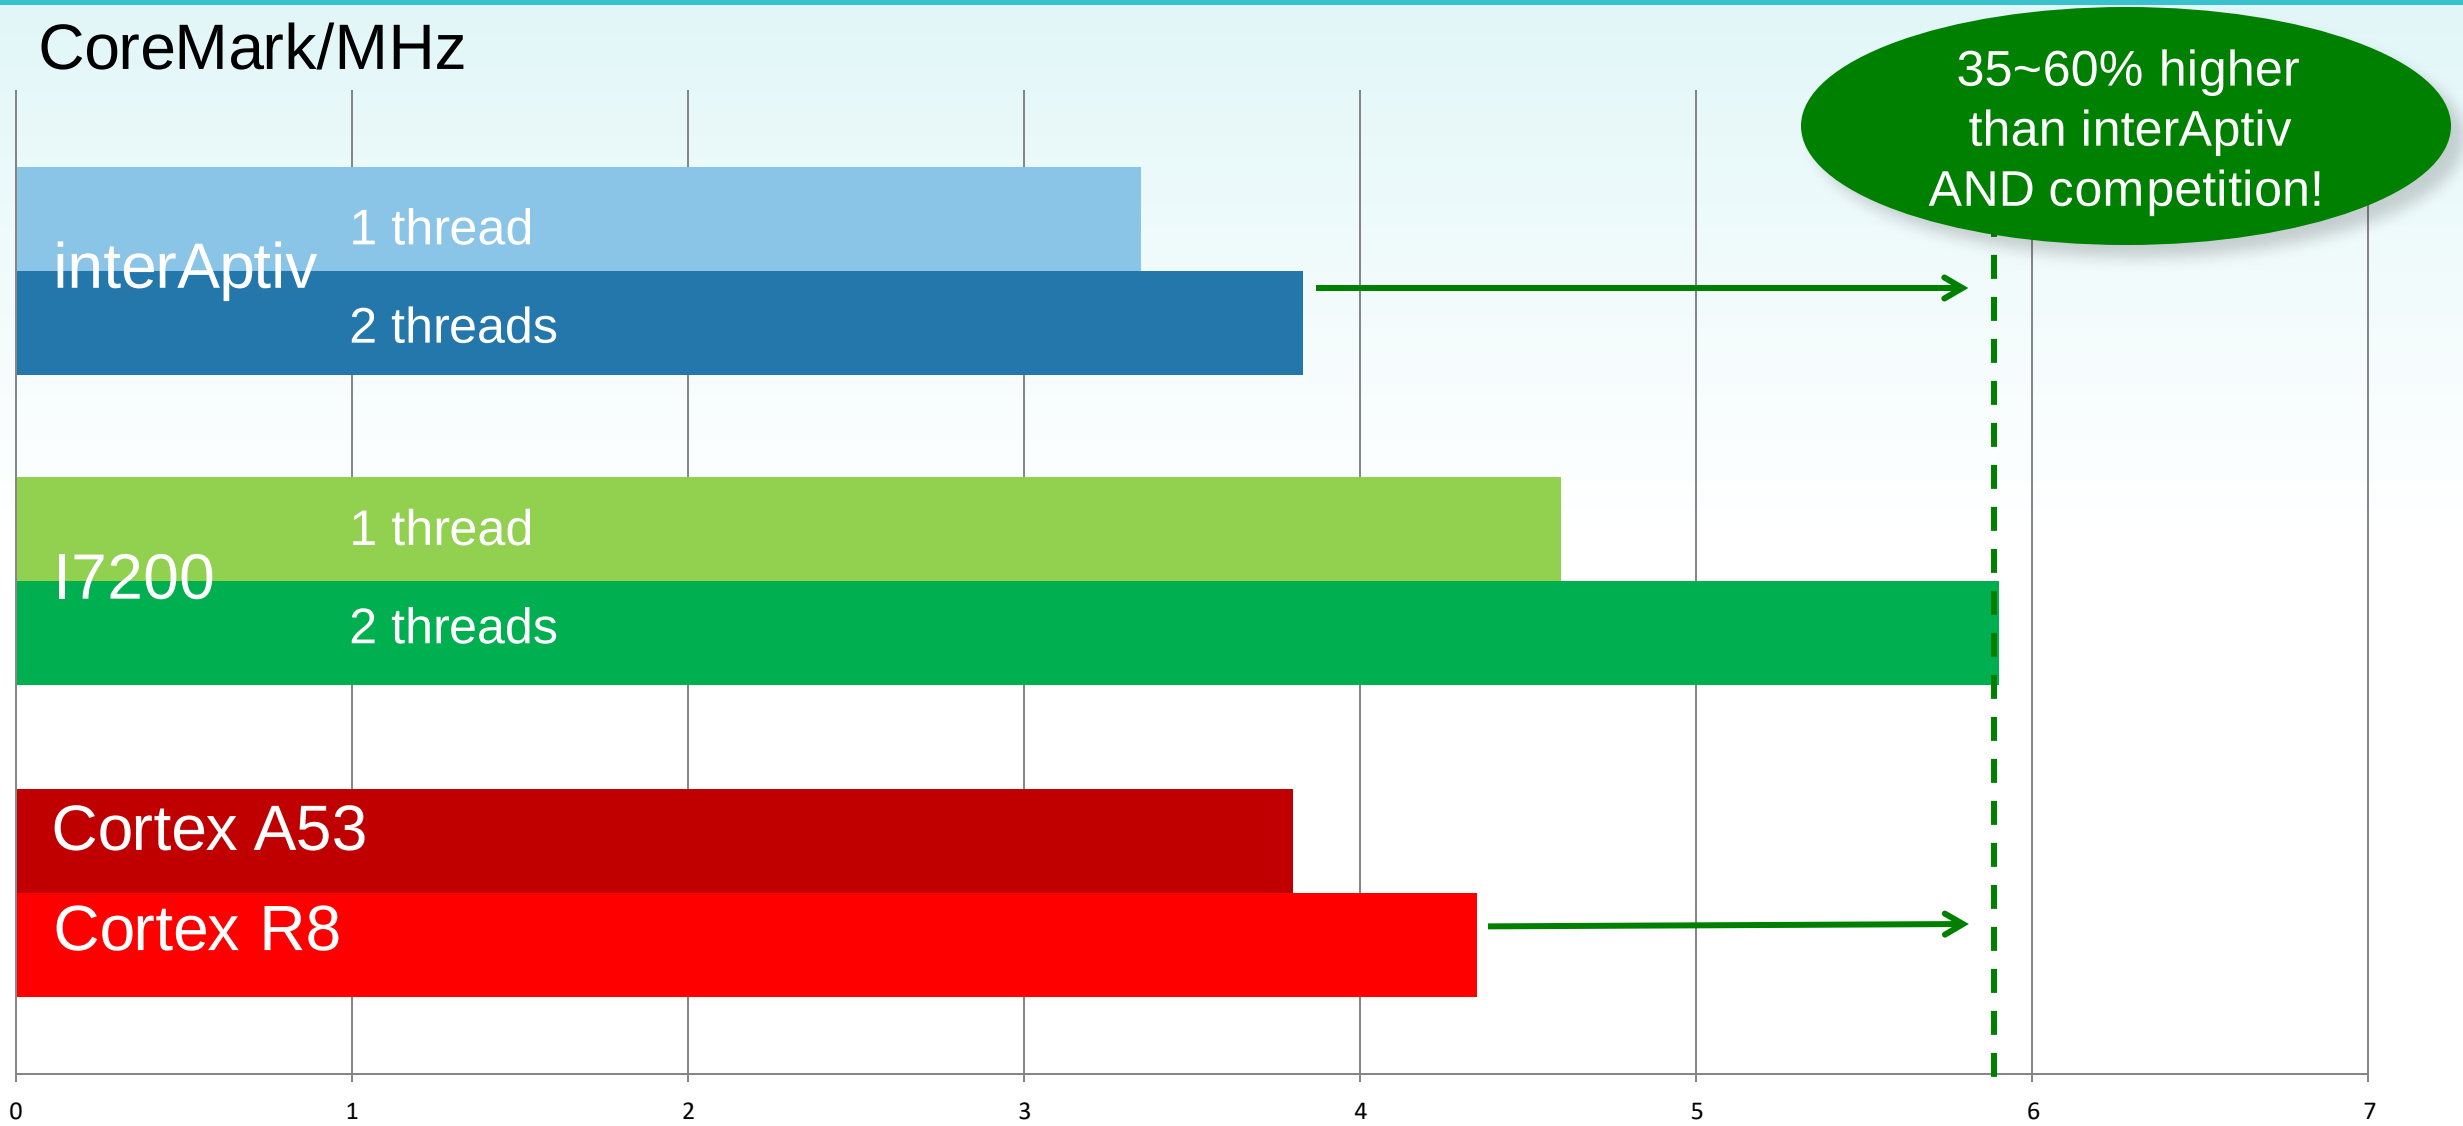

ARM scores reported by ARM for AArch32. Thumb2 scores likely 5-10% lower.
MIPS scores are MIPS32 for interAptiv, and nanoMIPS for I7200

WAVE COMPUTING

Application Segment 3
Data-center

Datacenter

Distributed Workload Processor

1. **Scalable - Multi-threading to Multi-core to Multi-cluster**
   - **Configurable - Heterogeneous inside & outside**
2. **Optimized for High-throughput data processing applications**
3. **Real-time, secure, deterministic and low latency**

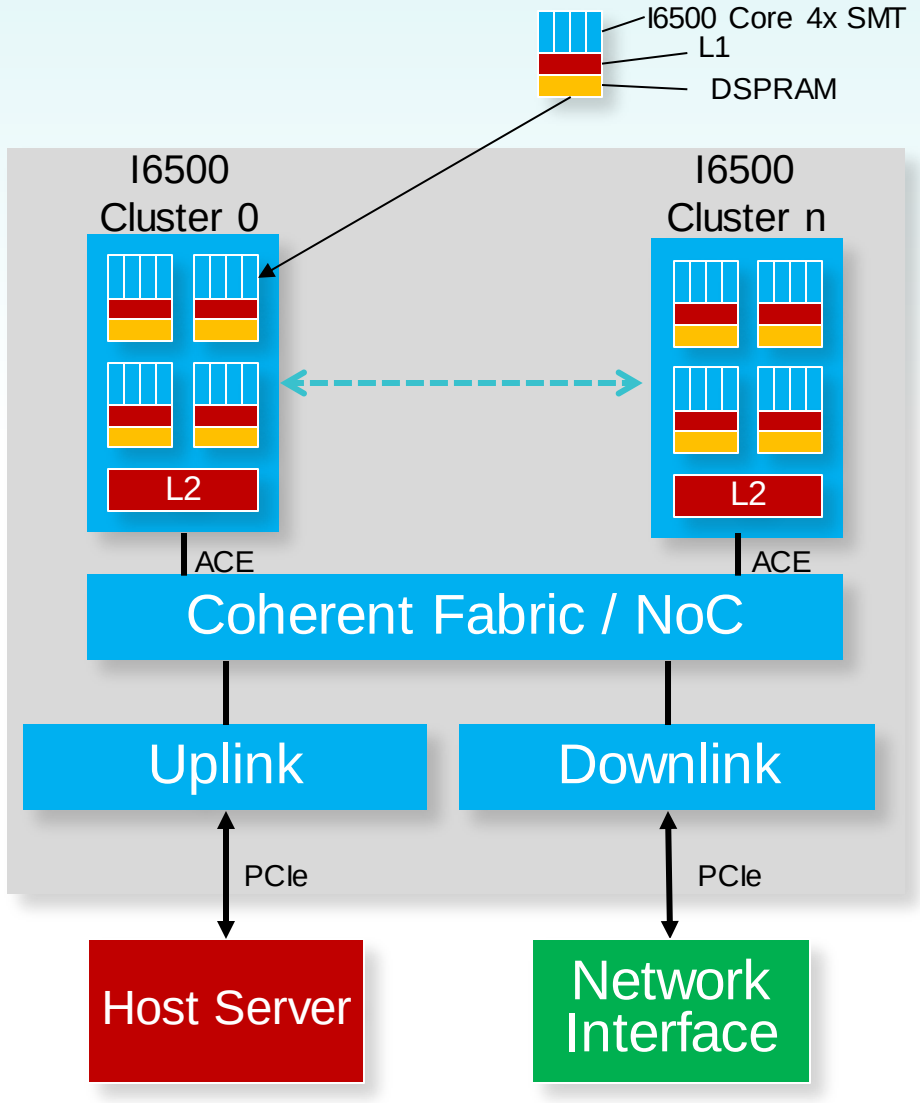

"The MIPS Simultaneous Multi-Threading architecture is an important technique to ensure that such workloads run efficiently as measured by the CPU's instructions per clock, or IPC. We have seen that this efficiency translates directly into a smaller area as well as lower power for silicon implementations based on MIPS." 2018

Pradeep Sindhu, CEO of Fungible

I6400 & 6500 – Why Multi Threading?

A powerful differentiator among CPU IP cores

1. **Why MT?**
   - A path to higher performance, and higher efficiency
   - 30%-60% higher performance for 10% per Thread increase
2. **Easy to use – programming model is same as multi-core**
   - A thread looks like a core to standard SMP OS
3. **Simultaneous / concurrent execution**
   - Zero Cycle overhead context switching

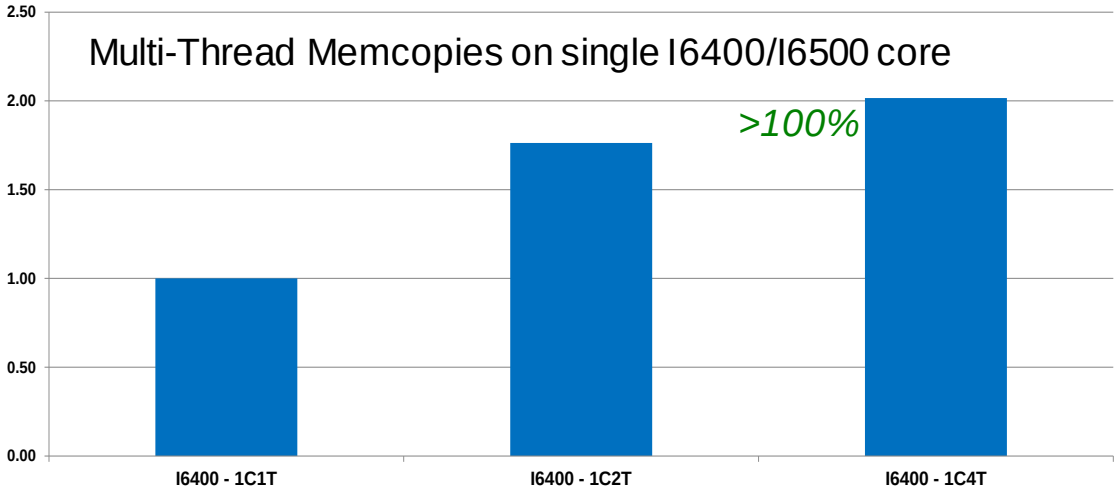

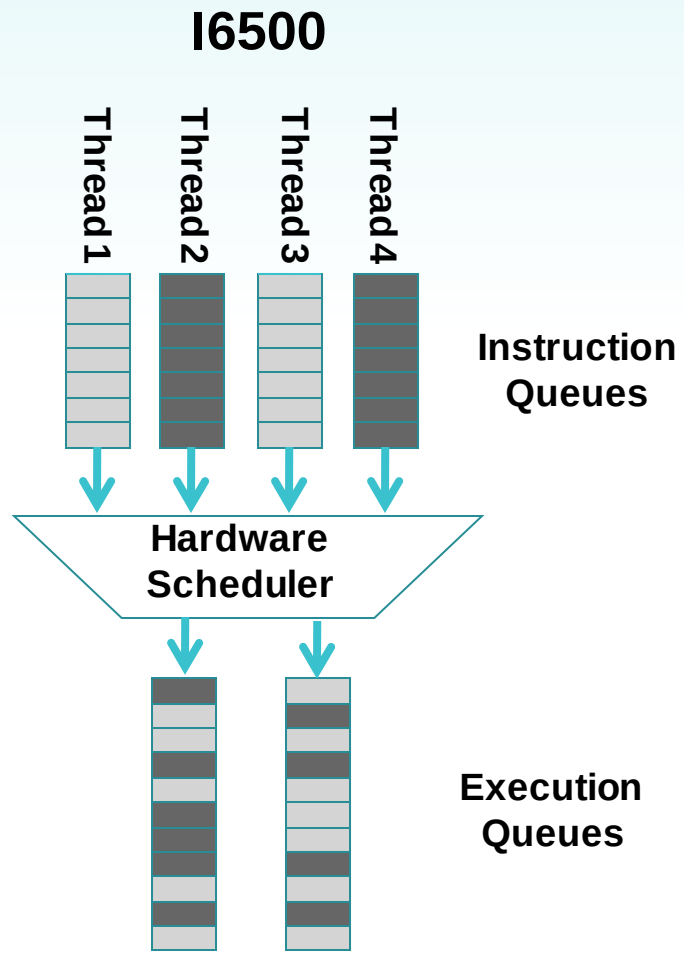

MIPS IP Cores – Mapping to ARM

	microAptiv	*M51xx*	M62xx	interAptiv	*I7200*	*I6500/-F*	*P5600*	*P6600*
MIPS Primary ISA	MIPS32 r5	MIPS32 r5	MIPS32 r6	MIPS32 r5	nanoMIPS32	MIPS64 r6	MIPS32 r5	MIPS64 r6
Virtual/Phys Addr Bits	32/32	32/32	32/32	32/32	32/32	48/48	32/40	48/40
FPU	✓ (UC version only)	✓	-	MT	-	MT w/SIMD	Hi Perf w/SIMD	Hi Perf w/SIMD
DSP/SIMD extensions	DSP ASE r2	DSP ASE r2	DSP ASE r2	DSP ASE r2	DSP ASE r2	MSA 128-bit	MSA 128-bit	MSA 128-bit
Virtualization	-	✓	-	-	-	✓	✓	✓
Small code size ISA	microMIPS32	microMIPS32	microMIPS32	MIPS16e2 ASE	nanoMIPS32	-	-	-
Multi-threading	-	-	-	2 VPE, 9 TC	3 VPE, 9 TC	4 VPE	-	-
SuperScalar	-	-	-	-	Dual-issue in order	Dual-issue in order	Multi-issue OoO	Multi-issue OoO
Pipeline stages	5	5	6	9	9	9	16	16
Relative Frequency*	0.6x	0.6x	0.75x	1x	0.95x	0.90x	1.10x	1.10x
SPRAMs (I/D/U)	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / -	✓ / ✓ / ✓	- / ✓ / -	- / - / -	- / - / -
L1 caches	✓	✓	✓	✓	✓	✓	✓	✓
L2 cache	-	-		✓	✓	✓	✓	✓
Coherent Multi-Core	-	-	-	Up to 4 cores	Up to 4 cores	Up to 6 cores, Up to 64 clusters	Up to 6 cores	Up to 6 cores
Native System Bus I/F	AHB-Lite	AHB-Lite	AXI	OCP 2 or AXI	AXI	AXI or ACE	AXI	AXI

Sample ARM Mapping

M3 M4 M23 M33 R52 R7 R8 A53 A57 A72

MIPS Open™

The New Standard in Open Use ISAs

Presented by Yuri Panchul, MIPS Open Technical Lead

MIPS Open Meetup in Moscow, April 15, 2019

Wave + MIPS – A Powerful AI + IP Combination

1981
MIPS Computer Systems Inc. founded

MIPS becomes IP licensing company

2017
MEDIATEK
Smartphone LTE modem

2019
Wave Establishes MIPS Open Advisory Board

1980 | 1990 | 2000 | 2010 | 2017 | 2018 | 2019

1981
MIPS architecture invented by team from Stanford

First MIPS product ships: R2000 uP

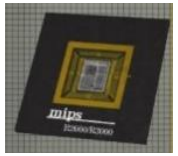

DTV/STB product ships:
NEC
SONY

Cable Modem product ships:
TOSHIBA
BROADCOM

Wi-Fi - Residential Gateway
ATHEROS
LANTIQ

Enterprise Networking
CAVIUM
PMC
NETLOGIC MICROSYSTEMS

Automotive MCUs & ADAS
MICROCHIP
MOBILEYE
MStar semiconductor

Wave Acquires MIPS
MIPS Open

Wave Launches MIPS Open

Joins AWS, Facebook, Google & Samsung on UC Berkeley's BAIR initiative

MIPS Open™: Accelerating Innovation @ the Edge

- Accelerates innovation for system-on-chip designs
- Expands adoption of MIPS RISC architecture & grows supporting ecosystem
- Dedicated Advisory Board prevents architectural fragmentation
- Delivers peace of mind & investment protection for existing MIPS customers

Decades of silicon-proven innovation


ADAS


TVs


Smartphones


Routers


Set-top boxes


Switches

Fueling Next-Gen Edge Designs


Autonomous vehicles


Visual analytics


IoT

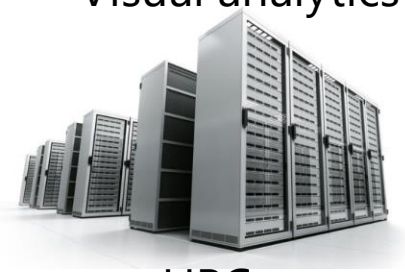
HPC

MIPS Open™ Benefits

Unified Approach

- Compatibility
- No Fragmentation
- Same tools, software
- Flexibility – UDI support

Proven Architecture

- MIPS architecture-based designs are shipping in billions of devices
- Shipping Products – wired modems, wireless modems – LTE/5G/Wi-Fi, IoT, automotive – ADAS, data centers, AI and others.

Established Ecosystem

- Well-established
- Tools, software, applications, boards
- Supported and used by customers/ partners

Matured Technology

- 30+ years of legacy of MIPS architecture
- Enabler to build the next trillion devices/ applications

MIPS Open is the World's First Commercial-ready, Silicon Proven, Feature-Rich Architecture Available for Open Use

MIPS Open™ Model

Training/Teaching/ Evaluation

Development Architecture to Core

Design & Build Core to SoC

Sell

MIPS Open FPGA

Open Source 32-/64-bit ISA DSP/SIMD/VZ/MT

Tools Software SDK

Verification Suite

Developed Cores

License Cores to Semi Vendors

Everyone – Student, University, Professional

Core Developer

Certification Partner

MIPS Open Certified Core

SOC#1 ..... SOC#n

Partner Tools Software SDK

Partner Design Services

Customer/ 3rd Party IP

MIPSOpen.com Now Live

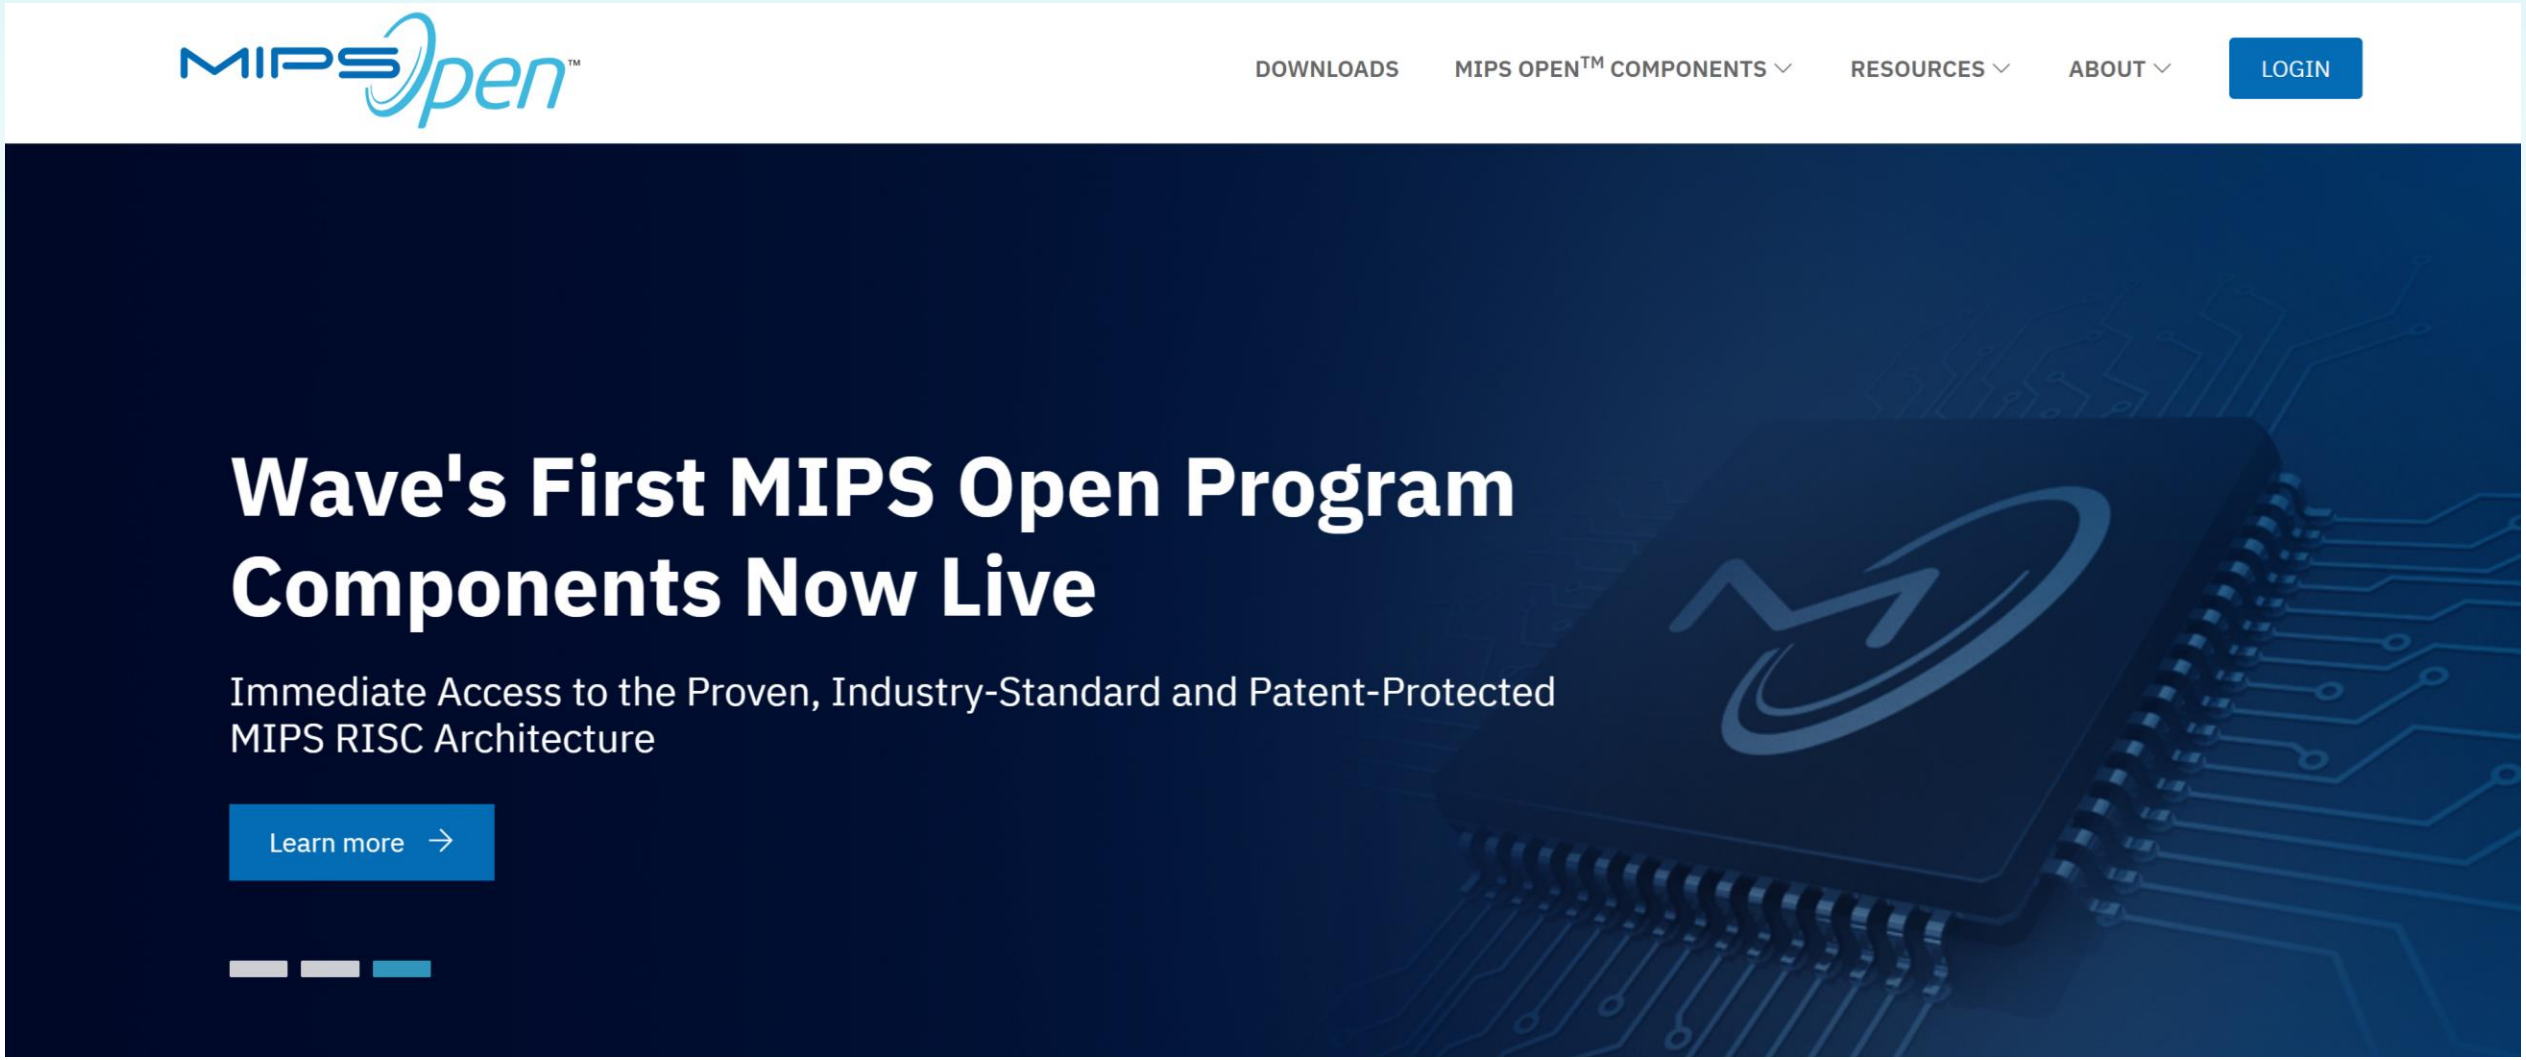

MIPS Open™ Components

MIPS Open ISA

- Latest R6 version of 32-bit/64-bit ISA
- Extensions such as Virtualization, Multi-threading, SIMD, DSP and microMIPS architectures

MIPS Open Tools

- IDE for Embedded RTOS and Linux Embedded Edition
- Enables MIPS software developers to build, debug, and deploy applications on MIPS-based hardware and software platforms

MIPS Open FPGA

- Getting Started Guide - provides MIPS Open FPGA system as a set of Verilog files
- Labs - includes 25 hands-on labs that guide users in exploring computer architecture & system-level design
- SOC - shows how to build a system-on-chip design based on MIPS Open FPGA that loads the open source Linux

MIPS Open Cores

- Low power, low footprint microAptiv cores
- Microprocessor(MPU)
- Microcontroller (MCU)
- Targeted for embedded applications

MIPS Open Components includes the latest MIPS R6 ISA, low power, low footprint microAptiv cores, microAptiv soft core targeted for FPGA and tools for software development.

MIPS Open™ Features

Open Use License

- MIPS Open use license includes:
  - Latest MIPS R6 architecture documents
  - microAptiv cores
  - Tools
  - microAptiv softcore targeted for FPGA

Certification

- Access to certification services to core developers
- Verification and Certification of cores developed using MIPS R6 ISA.

Patent License

- Right and license under R6 architecture patents to design, build and sell cores
- Use of the "MIPS Certified" trademark logo for certified cores

Comprehensive Package

- Complete package
- Instruction set, cores, tools for the community to accelerate innovation at the edge

MIPS Open provides a comprehensive open use package and tools to enable fast time to market for certified cores.

Evolving the MIPS Ecosystem for AI at the Edge

Building on an established ecosystem…

Mixed Signal

Memory

Growing new markets with AI

WAVE COMPUTING

…creating an AI solution stack that enables end-to-end business models

Today → *Tomorrow*

MIPS Open™ Advisory Committee Membership Levels

Individual Membership

Entry level

FREE
per year

- Allows participation in all working groups and access to all working group materials.
- Represent individual, academic and non profits interests.

Silver Membership

Entry level of corporate membership.

$ 10,000
per year

- Allows participation in all working groups and access to all working group materials.
- Right to vote on working group initiatives.
- Can be appointed to lead a working group. If so appointed, automatically become a member of the technical steering committee.
- Can be nominated to run for election to the board of directors as a Silver class representative.
- Vote as a Silver class for representation on the board of directors.

Gold Membership

Middle tier of membership.

$ 50,000
per year

- Allows participation in all working groups and access to all working group materials.
- Right to vote on working group initiatives.
- Can be appointed to lead a working group. If so appointed, automatically become a member of the technical steering committee.
- Can be nominated to run for election to the board of directors as a Gold class representative.
- Vote as a Gold class for representation on the board of directors.Allows participation in all working groups and access to all working group materials.

Platinum Membership

Decision-making level of membership

$ 100,000
per year

- Leadership and decision-making level of membership.
- Automatically appointed to the board of directors and the technical steering committee.
- Set direction, approve budgets and projects, appoint working group leaders, create new working groups, and supervise foundation staff.
- Have access to all foundation meetings and materials.
- Allows participation in all working groups and access to all working group materials.
- Right to vote on working group initiatives.Leadership and decision-making level of membership.

MIPS Open™ Timeline

Q2'19

'Open Use' microAptiv UC/UP Cores

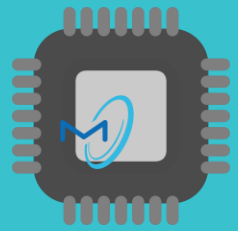

March 28th, 2019

Feb 19th, 2019

Dec 17th , 2018

Thank You