



INTRODUCING “PARKER”

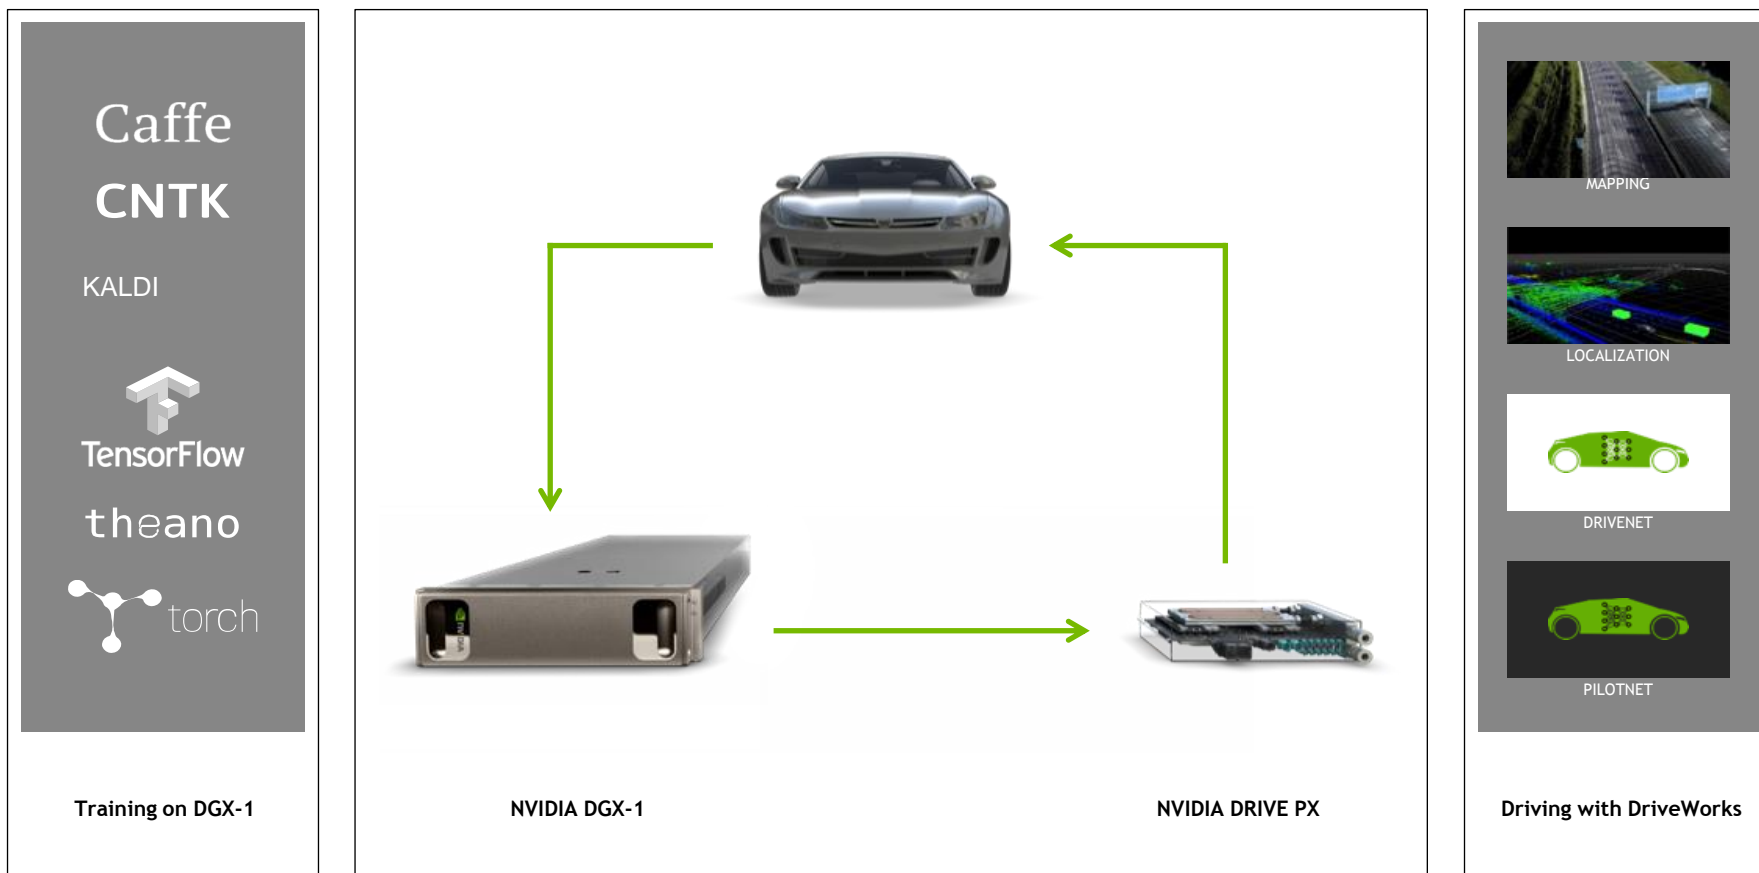
Next-Generation Tegra System-On-Chip

Andi Skende | Distinguished Engineer, “Parker” Lead Architect

“PARKER” DESIGN OBJECTIVES

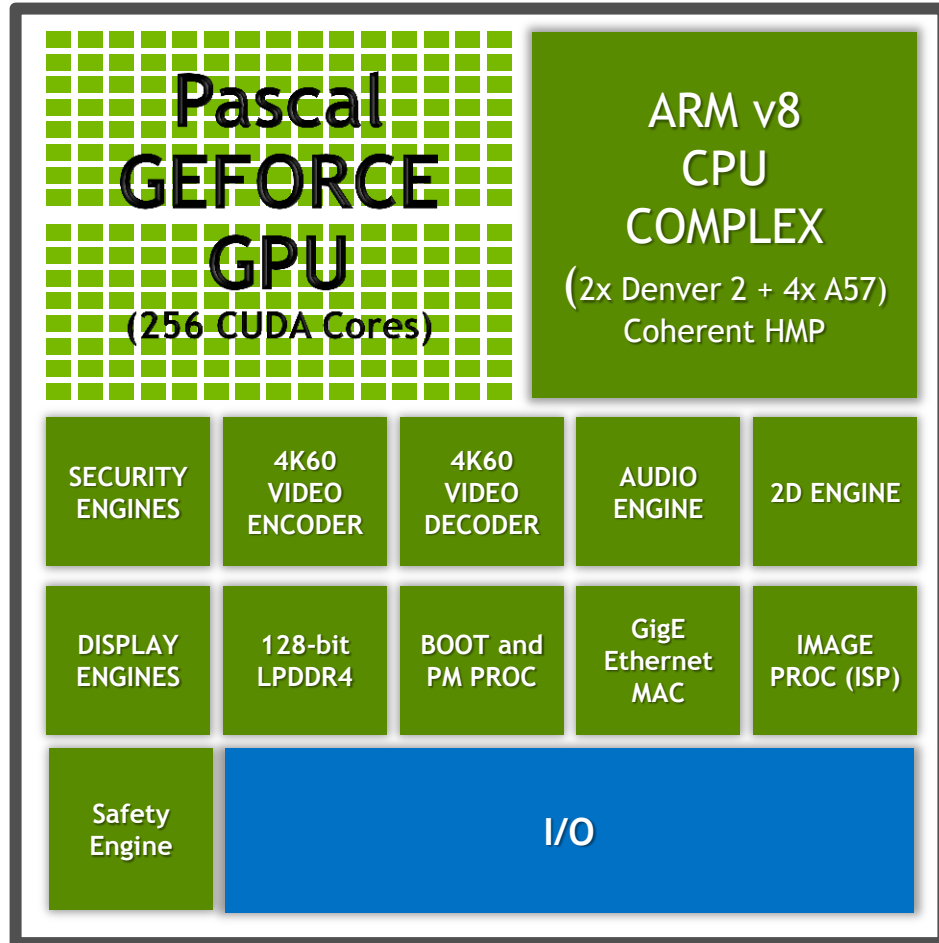
- Deliver the most advanced processor for premium automotive market
 - A deep learning computing platform for autonomous machines
 - Scalable architecture: autopilot to self-driving
 - A single virtualized machine

NVIDIA DRIVE END TO END DEEP LEARNING PLATFORM



“PARKER”

NEXT-GENERATION SYSTEM-ON-CHIP



- NVIDIA’s next-generation Pascal graphics architecture
- NVIDIA’s next-generation ARM 64b Denver 2 CPU
- Functional safety for automotive applications
- Hardware-enabled virtualization architecture
- Improvements to SoC architecture to enable modularity and ASIC development efficiency
- Industry-leading 16nm FF process

TEGRA KEY FEATURE EVOLUTION

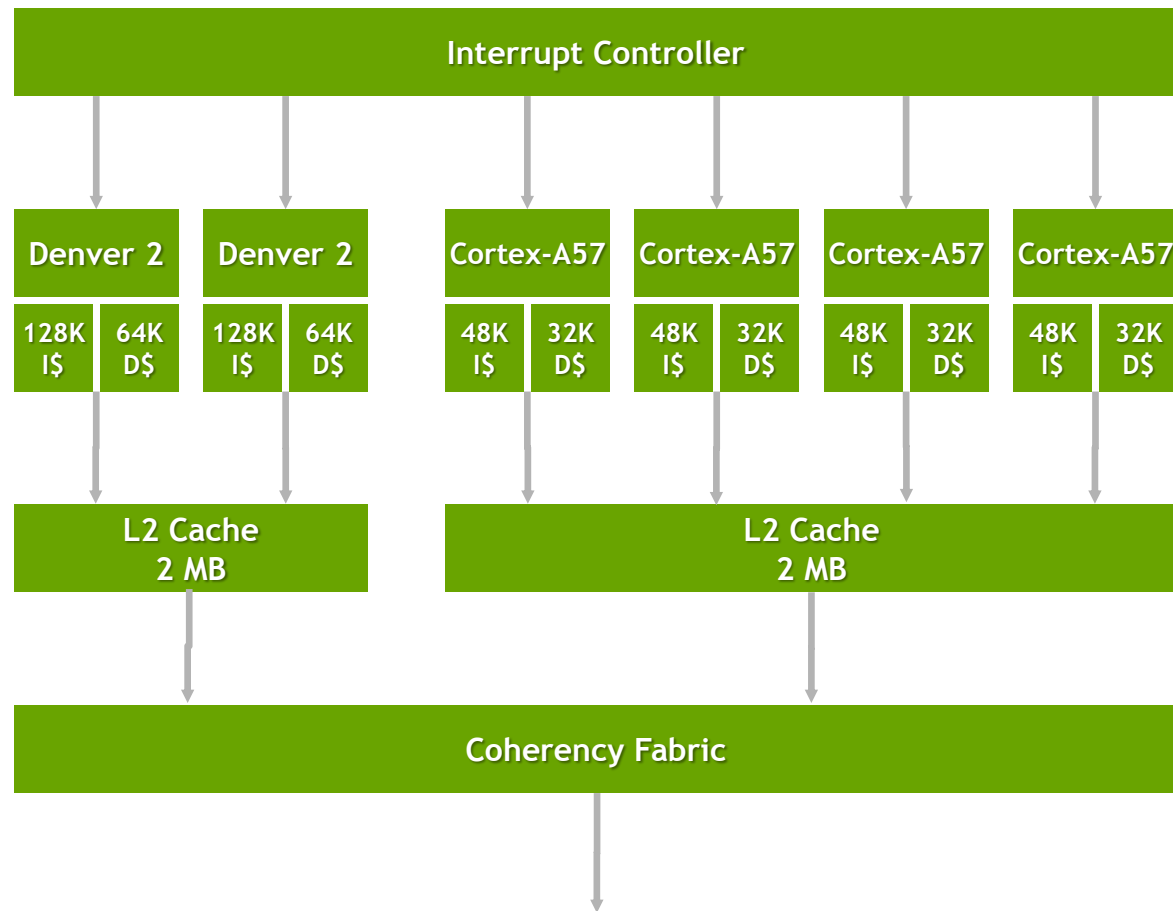
	TK1	TX1	“PARKER”
GPU	Kepler, 192 CUDA cores	Maxwell, 256 CUDA cores	Pascal, 256 CUDA cores 1.5TFlops (fp16)
CPU	4+1 A15, 2MB+512K L2 ARM v7 32b Or 2 Denver 1, 2MB L2 64b	4x A57 2MB L2 + 4x A53 512KB L2 ARM v8 64b	2x Denver 2 2MB L2 + 4x A57 2MB L2 ARM v8 64b Coherent HMP Architecture
Camera	4 cameras	6 cameras	Auto HDR 12 cameras
Memory	64b LPDDR2/3, DDR3L 15 GB/s (LP3, DDR3L)	64b LPDDR4, 25GB/s	128b LPDDR4, ~50 GB/s, ECC
Display	Dual Pipeline 4K@30fps 24bpp	Dual Pipeline 4K@60fps	Triple Pipeline 4K@60fps

TEGRA KEY FEATURE EVOLUTION

	TK1	TX1	“PARKER”
Video	2160P30 decode, 2160P30 encode	2160P60 decode, 2160P30 encode	2160P60 decode, 2160P60 encode
Storage	e-MMC 4.51, SATA	e-MMC 5.1, SATA	e-MMC 5.2, SATA
Auto Features	NA	QSPI	Ethernet-AVB, Dual CAN, QSPI
Resiliency / Safety Features	NA	NA	Automotive rated SoC Extensive set of resiliency features On-die Safety Manager Engine
Virtualization	NA	NA	HW-assisted Virtualization
Process	TSMC 28HPM	TSMC 20SOC	TSMC 16FF

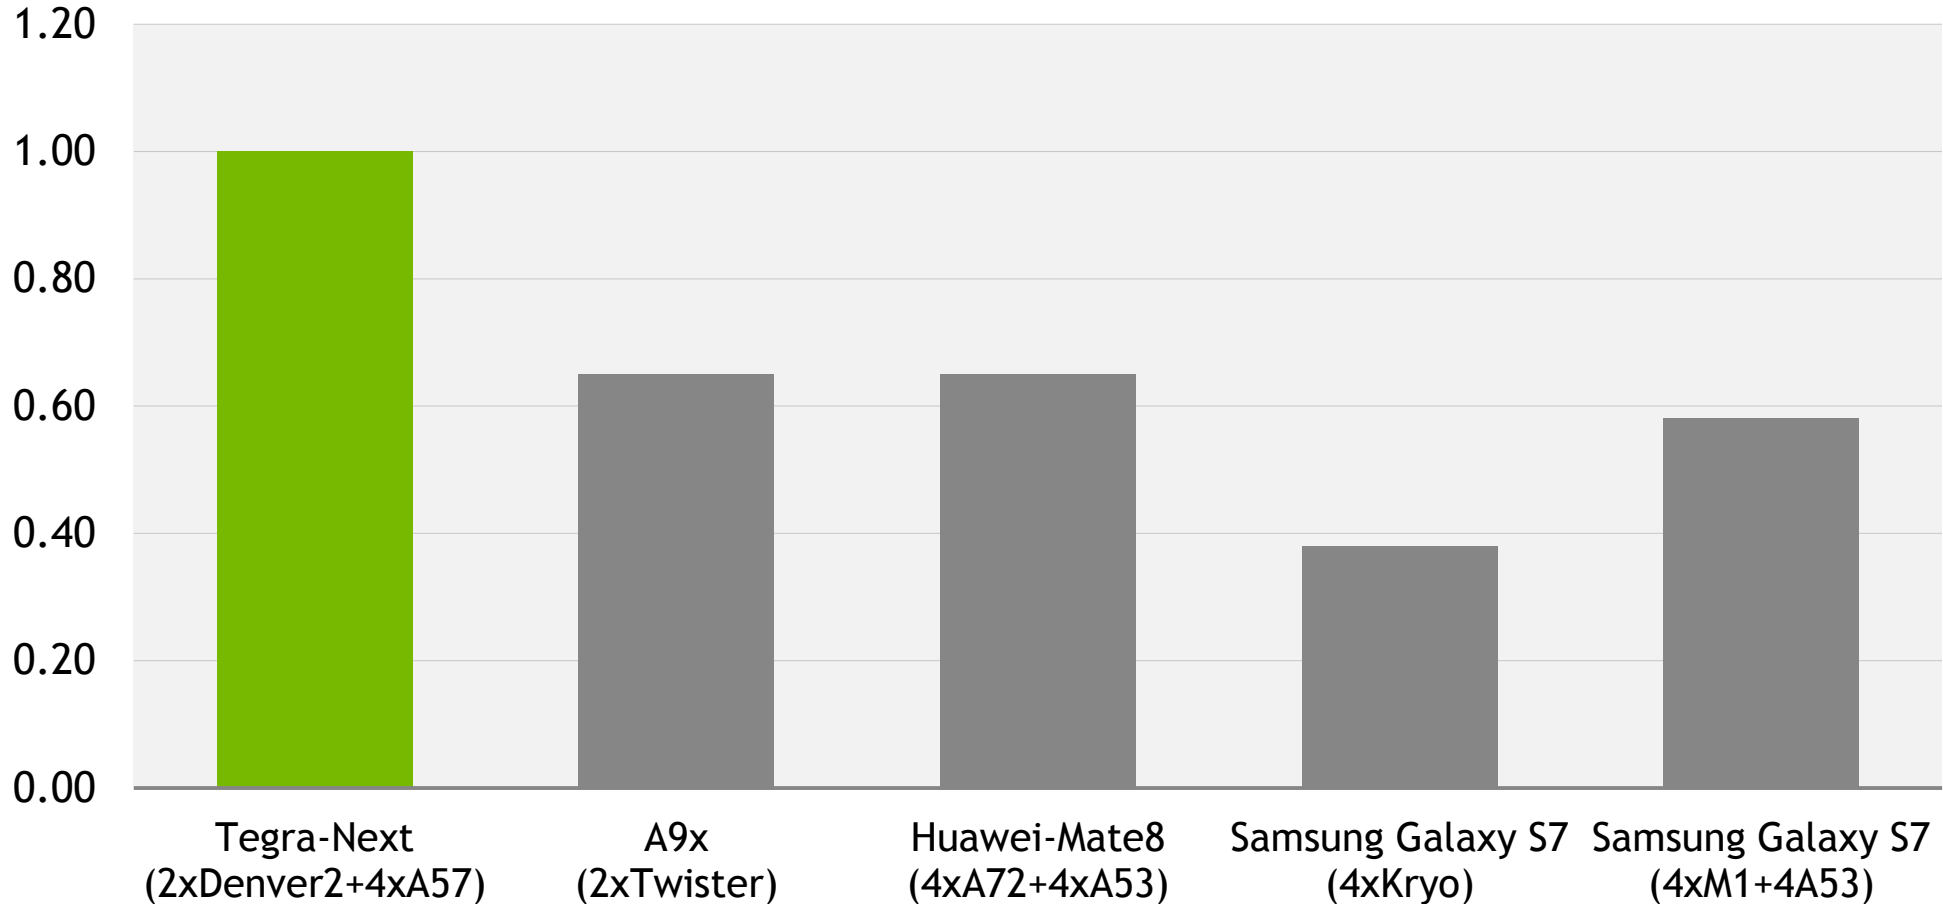
“PARKER” CPU COMPLEX

- **2x Denver2 + 4x Cortex-A57**
- **Fully Coherent HMP system**
 - Proprietary Coherent Interconnect
- **ARM V8 64-bit**
- **Highest performance ARM CPU**
 - 2nd generation Denver core
 - Significant Perf/W improvements (~30%)
- **Dynamic Code Optimization**
 - Optimize once, use many times
- **7-wide superscalar**
- **Low power retention states**



“PARKER” CPU PERFORMANCE

SpecInt2K-Rate Relative to “Parker”



COHERENT HETEROGENEOUS MULTIPROCESSOR

	“PARKER” - COHERENT HMP CPU ARCHITECTURE
Heterogeneous	<ul style="list-style-type: none">• Schedule the task on the right CPU core• Move task/thread to the right core as compute needs change• Maximize peak performance and responsiveness
Coherent	<ul style="list-style-type: none">• No software overhead for cache maintenance as the thread moves across clusters
MultiProcessor (Big + Super)	<ul style="list-style-type: none">• Great single thread performance• Maximize aggregate performance• Sufficient thread count for automotive and gaming applications

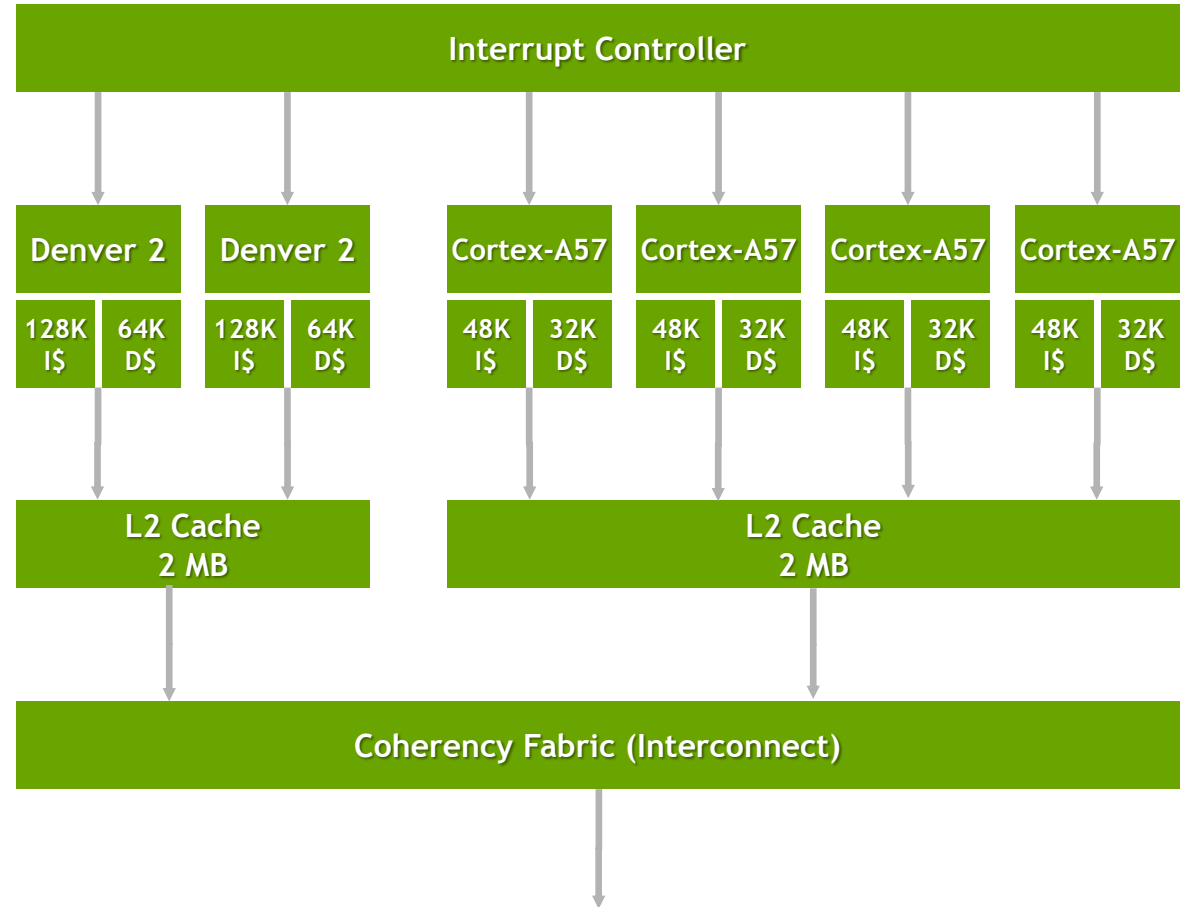
COHERENT HMP - DESIGN CHALLENGES

- **Coherency Interconnect**

- Enable coherency across two different micro-architectures (Denver vs. A57)

- **Power Management**

- Unified power state and power management architecture across Denver and A57 cores



VIRTUALIZATION ON “PARKER”

HW-enabled CPU, GPU and SoC Virtualization



- HW-enabled virtualization
- Up to 8 virtual environments (VMs)
- Virtualization enables integration
 - Example: IVI + IC in one ECU
- Each VM can control its own display pipeline
- Supported by highest quality virtualization software from NVIDIA

VIRTUALIZATION ON “PARKER”

Hardware Features

- **ARM CPU virtualization extensions**
 - HYP privilege mode for Hypervisor execution
 - vCPU and vGIC state
 - Virtual CPU timer
 - Two-stage MMU translation for access isolation
 - Instruction trapping (e.g. “WFI”)
- **System MMU**
 - Two-stage address translation
 - Isolation and protection for memory accesses
- **Hardware checks**
 - Configuration apertures aligned to 64KB
 - Detection of illegal accesses
- **GPU**
 - Multiple physical channels directly accessible to VMs
 - GPU MMU for isolating many channel contexts
 - Hardware assisted context switching
 - Fine granularity pre-emption for Graphics
 - CUDA compute context preemption at instruction-level boundary
- **Audio Processing Engine**
 - Per VM Audio DMA channel assignment and isolation
- **General Purpose DMA**
 - Per VM DMA channel assignment and isolation
- **Virtualization aware IO controllers**
 - Ethernet-AVB, PCI-E, etc.

“PARKER” - AN AUTOMOTIVE SoC

• Auto Resiliency Features

- ISO-26262 compliant
- Extensive support for error detection and correction
- Safety Engine
 - Dual lock-step RT processor
 - ▶ Supports error/fault management and reporting
 - ▶ Executes safety diagnostics and recovery libraries
- In-line DRAM ECC
 - Errors reported to Safety Engine
 - HW-assisted patrol scrubbing and DRAM page blacklisting
- ECC / parity protection for key on-die memories

• Automotive Centric IO(s)

- CAN Interface
 - Dual CAN interface
 - Always-ON domain for fast wake-up
 - ▶ Conforms to ISO11898-1
 - ▶ Time-triggered interface between CAN controllers
- Ethernet Audio Video Bus
 - Gigabit Ethernet MAC
 - RGMII interface to external PHY or switch
 - ▶ Real time bandwidth reservation

“PARKER” AND DRIVE PX 2

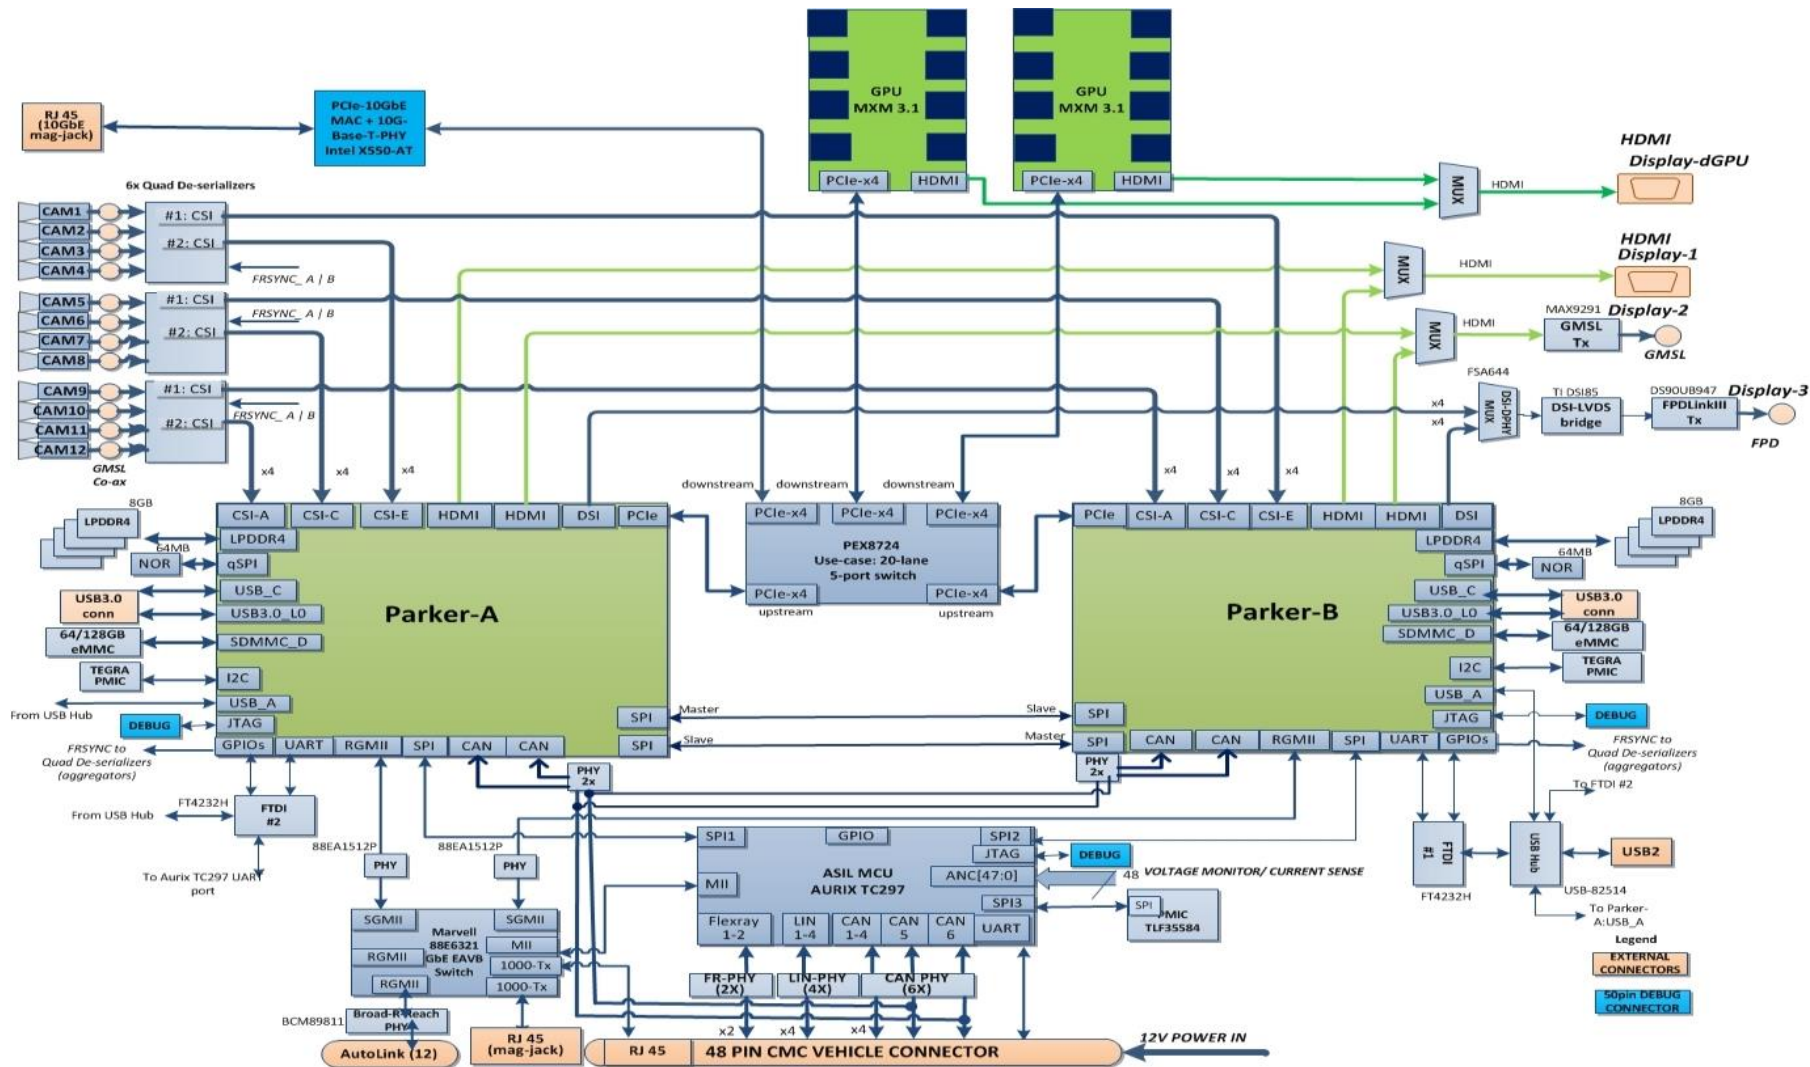
12 CPU cores | Pascal GPUs | 8 TFLOPS | 24 DL TOPS



- Scalable
 - Up to two “Parker” chips + two dGPU chips
- Developed as SEooC (Safety-Element-out-of-Context)
 - Provides flexibility to the developers in supporting their specific use cases and safety goals
- Deployed to many OEM and Tier1 partners
 - Supports development of Autonomous Driving solutions

World’s First AI Supercomputer for Self-Driving Cars

“PARKER” AND DRIVE PX 2



SUMMARY - TEGRA “PARKER”

- Improved perf/w with Denver 2 CPU core
- Exceptional aggregate (big + super) CPU performance
- Easy of programming through HMP architecture
- Best-in-class integrated GPU core based on the latest Pascal architecture
- A flexible automotive SoC
 - autopilot to self-driving use cases
- The main building component of NVIDIA DRIVE PX 2 platform



nVIDIA.

