



# Increasing SerDes, Memory Bandwidth and Security for Networking, 5G, AI, HPC and IOT

David Kuo


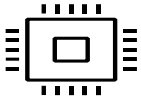
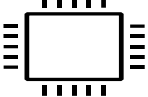

Rambus Greater China

Sep. 12, 2019



# Rambus at a Glance

## Rambus Offerings

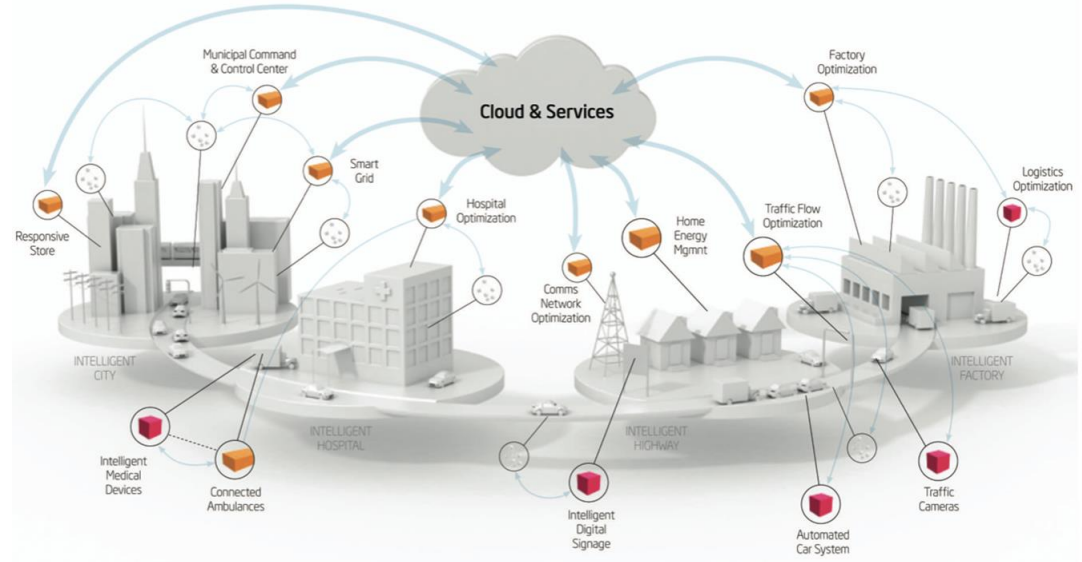
Architecture Licenses		High-speed IO & DPA Countermeasures
IP Cores		Memory & SerDes PHYs; Secure Cores
Chips		Memory Buffers for DIMM modules
Key Management		Secure Supply Chain Provisioning



# Our Increasingly Connected World Continues to Evolve

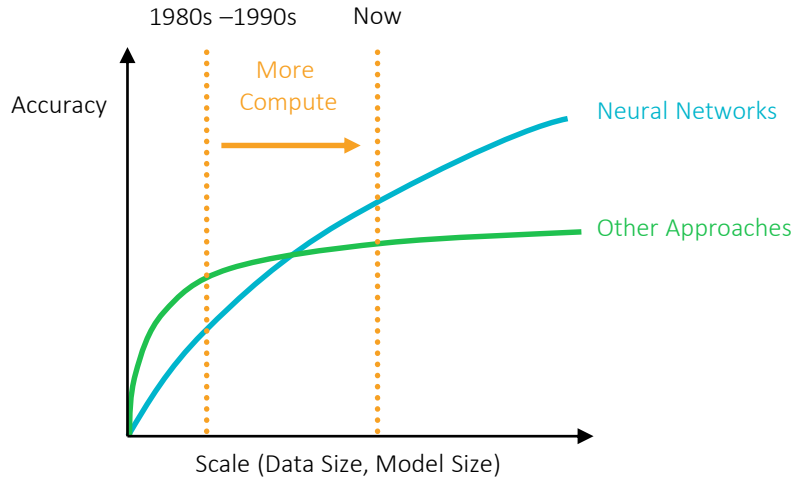
Advances in mobile devices, cloud computing, artificial intelligence, and connectivity are driving the growth of The Internet of Things

- Data and insights are increasingly valuable, security becoming more important
- Memory, link, storage performance must continue to increase
- Compute and I/O power efficiency must also continue to improve




Modern applications are changing the way we process information

# Faster Compute + Big Data Enabling Explosive Growth in AI



Source: Adapted from Jeff Dean, "Recent Advances in Artificial Intelligence and the Implications for Computer System Design," HotChips 29 Keynote, August 2017

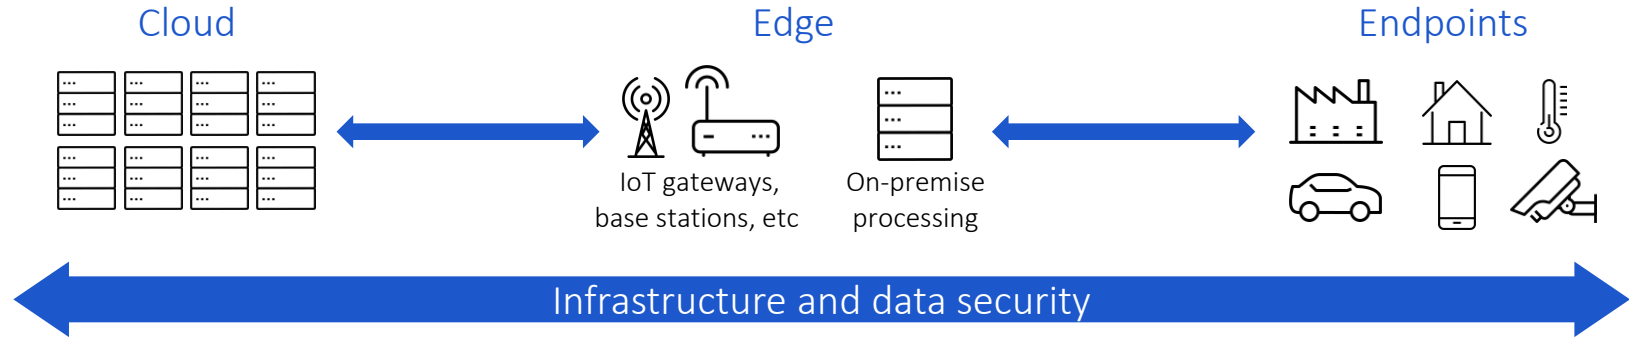
### Annual Size of the Global Datasphere



Source: Adapted from Data Age 2025, sponsored by Seagate with data from IDC Global DataSphere, Nov 2018

- 2019 Turing Award: LeCun, Hinton, and Bengio for key developments in Neural Networks
- Key challenges: Moore's Law ending, energy efficiency growing in importance

# AI Needed Across the Evolving Internet




- Analyze broad behaviors, train neural networks for cloud, edge, endpoints
- Need **high performance, low power training**, also need **low power inference**
- Need highest performance memories (on-chip, HBM, GDDR), power efficiency also critical

- **5G** brings new opportunities for **training** and **inference**
- Analyze endpoint data, provide broader actions and summaries
- Need **high performance, low power training** and **low power inference**
- Range of memories needed: on-chip memory, HBM, GDDR, DDR, LPDDR

- Local processing improves security, allows communication of a smaller amount of processed data instead of raw data
- Need **low power inference**
- Typically use on-chip memory or low-power mobile memory

# We Are At The Beginning Of A “Tectonic Shift In Computing”




- Parallelism
- AI & Machine learning
- Heterogenous architectures
- Low cost storage
- 5G /WiFi / BLE Wireless Connectivity

Order of magnitude volume Increase

# Memory and Networks Limiting System Performance


Computer systems are falling further out of balance



Source: [hpcwire.com/2016/11/07/mccalpin-traces-hpc-system-balance-trends](https://hpcwire.com/2016/11/07/mccalpin-traces-hpc-system-balance-trends)

- Compute improving much faster than memory and networks
- Memory and network latency and bandwidth falling behind processor performance at an alarming rate
- Memories and networks are becoming even larger bottlenecks

# Ethernet Roadmap – Link Speed

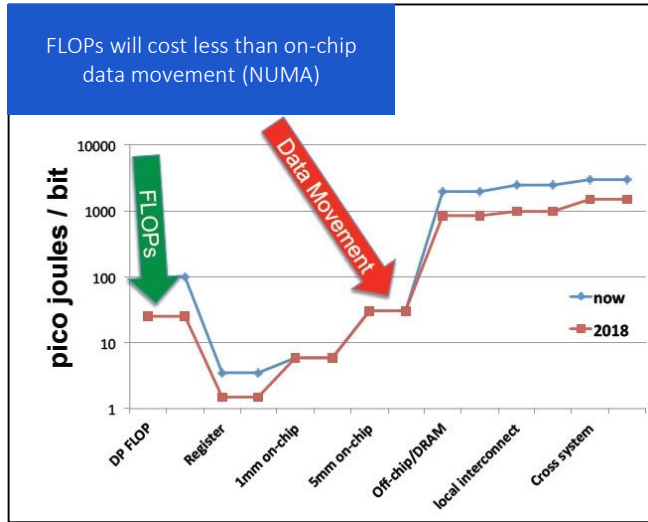


Source: 2019 Ethernet Alliance Roadmap  
[ethernetalliance.org/wp-content/uploads/2019/02/EthernetRoadmap-2019-Side2.jpg](https://ethernetalliance.org/wp-content/uploads/2019/02/EthernetRoadmap-2019-Side2.jpg)



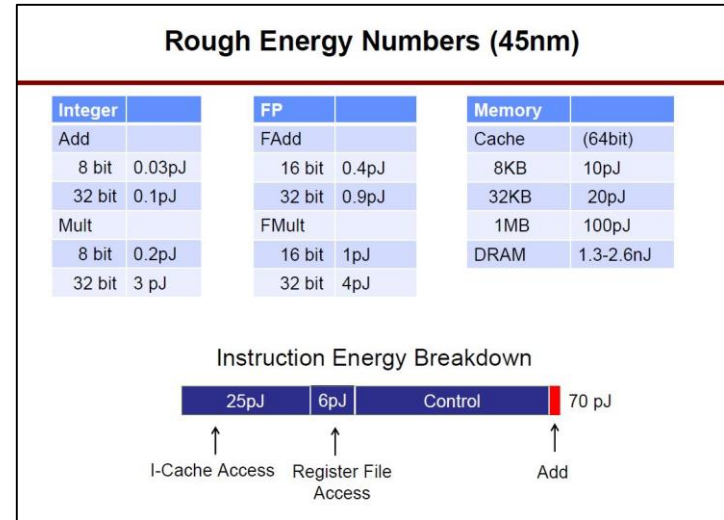
# Power Dominated by Data Access and Data Movement

2013: Computation Will Take Less Energy Than On-Chip Data Movement by 2018



Source: Horst Simon, "Why we need Exascale and why we won't get there by 2020," 2013

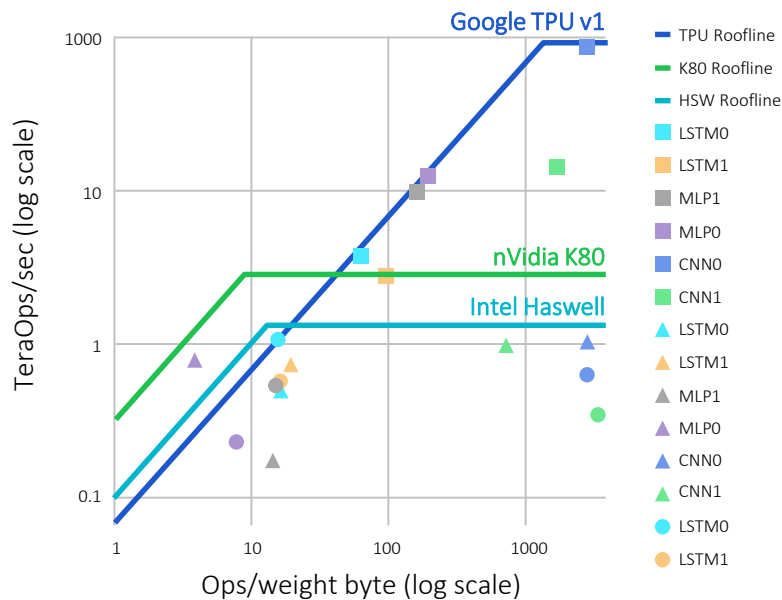
2014: Computation Energy Extremely Small Compared to Data Access and Movement



Source: Mark Horowitz, ISSCC "Computing's energy problem (and what we can do about it)," 2014.

There is a similar imbalance between compute time and data movement time, need to minimize data movement!

# AI Accelerators Need Memory Bandwidth



□ = Google TPU v1

△ = nVidia K80

○ = Intel Haswell

- Inference on older, general purpose hardware (Haswell, K80) limited by compute and memory bandwidth
- Inference on newer silicon (Google TPU v1) built for AI processing largely limited by memory bandwidth

Memory bandwidth is a critical resource for AI applications

Adapted from N. Jouppi, et.al., "In-Datacenter Performance Analysis of a Tensor Processing Unit™,"  
[arxiv.org/ftp/arxiv/papers/1704/1704.04760.pdf](https://arxiv.org/ftp/arxiv/papers/1704/1704.04760.pdf)

*Rambus*

## Memory PHY

David Kuo

Rambus Greater China

Sep. 12, 2019




R



# Common Memory Systems for AI Applications

## On-Chip Memory




Highest Bandwidth  
and Power Efficiency


Apple A11  
Bionic Processor



Graphcore  
IPU



## HBM




Very High Bandwidth  
and Density

nVidia  
Tesla V100



AMD Radeon  
RX Vega 56



## GDDR




Good tradeoff between bandwidth,  
power efficiency, cost, and reliability

nVidia GeForce  
RTX 2080Ti



AMD Radeon  
RX580



Multiple options suited to different needs

# Emerging Applications: Driving Compute/Memory Evolution

*Memory Bandwidth is a Critical Bottleneck!*

## Applications

- **Networking:** Wireline and Next Generation Wireless Infrastructure
- **Automotive:** ADAS and full autonomous >100GB/s
- **AI: Machine Learning and Deep Learning:** Advanced neural networks


## Custom silicon and new system architectures

- Emerging applications reaching the limits of current hardware architectures
- Increasing reliance on accelerators and specialized silicon such as neural networks
- Memory IP is critical for addressing a key bottleneck in these systems


## Advanced Driver Assistance Systems



## Machine Learning / Neural Networks



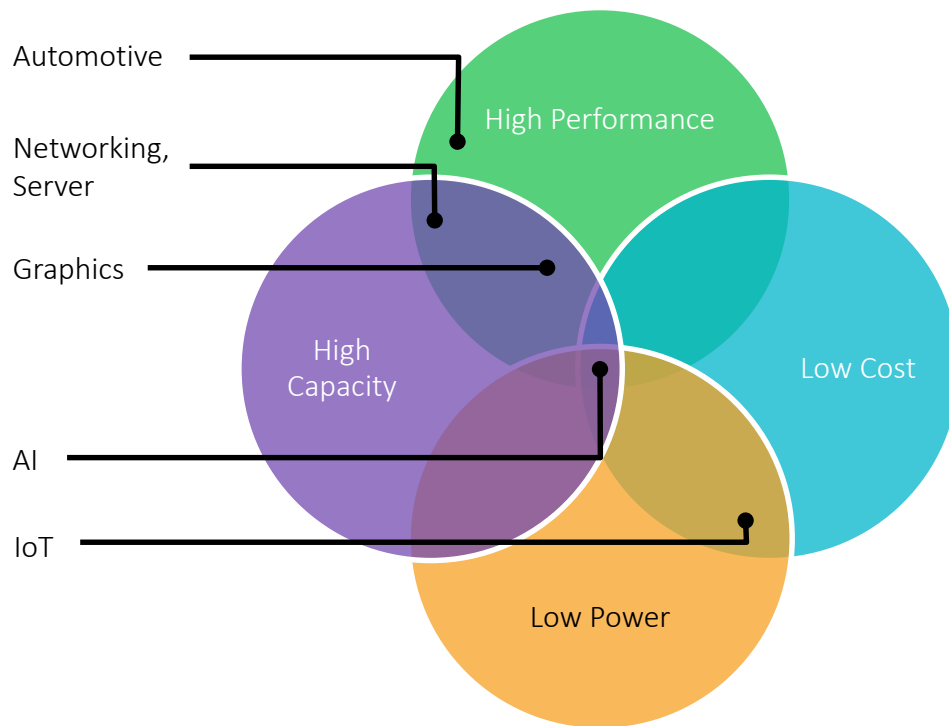
## Blockchain / Cryptocurrency Mining



# Key Memory System Metrics

- Cost
- Power
- Capacity
- Performance

Trade off is required based on application requirements while selecting memory solution !!



# Memory Comparison

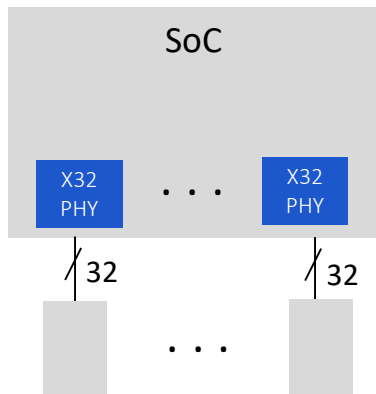
Parameter	GDDR6	HBM2	DDR4	LPDDR4X
Bandwidth (Gbps)	High (512)	Highest (2000)	Medium (200)	Low-Medium (136)
Data Rate (Gbps)	16	2	3.2	4.266
Interface width (bits)	32	1024	64	32
Board Area / System Design	Medium / Medium	Small / Complex	Large / Easy	Large / Medium
Efficiency (mW/Gbps)	Moderate (10)	Highest (2)	Moderate (10)	High (3)
Cost (\$)	Medium	High	Low	Medium
Reliability/Yield	Good	Moderate	Good	Good

# Memory System Comparison: LPDDR4 vs. GDDR6 vs. HBM2

## Memory System Comparison: LPDDR4 vs. GDDR6 vs. HBM2

### LPDDR4

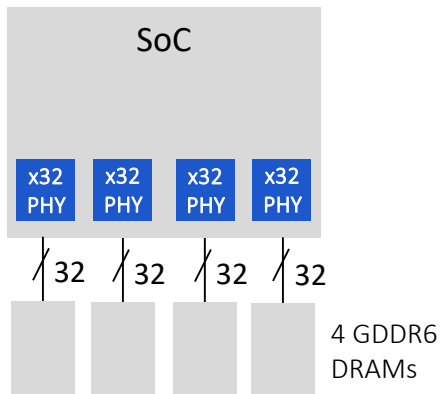
256GB/s Memory System



- 15 x 32 DDR4 DRAMs @4.266 Gbps
- Standard PCB manufacturing
- Large board and die area
- Challenging to implement: long channel (terminated)

### GDDR6


256GB/s Memory System



- 4 GDDR6 PHYs/DRAM @16Gbps
- Standard PCB manufacturing
- Higher board area
- Lower overall system cost
- Higher power

### HBM2

256GB/s Memory System



- 1 HBM2 PHYs/DRAM @2Gbps
- 2.5D manufacturing
- Lower board area
- Higher cost
- Lower power


# Research Enabling Rambus' GDDR6 PHY

Leading-Edge Technology Development Leveraging Advanced System-Level Modeling and Test Chip Development Capabilities

## 2007 Test Chip/System

- Innovations, building blocks for future TB/s memory systems
- Silicon and system demo of signaling at 16+Gbps


### Silicon and System Demo




Package Options



System Validation Board



Controller equalized  
16 Gbps TX eye



16 Gbps Read  
Bathtub Curve


## 2010 Test Chip/ System

- Multi-modal controller PHY with 20+Gbps differential signaling, GDDR5, and DDR3
- 12.8Gbps single-ended signaling with high reliability


### Silicon and System Demo




System Board




Single-ended 12.8 Gbps



XDR2 20 Gbps



DDR3 2.4 Gbps



GDDR5 6.4 Gbps


## 2019 Test Chip/ System

- Validation Vehicle for GDDR6 memory interface and system
- TSMC 7nm FinFET process
- Hard macro solution designed for easy integration into ASICs
- LabStation™ development environment for quick system bring-up, debug, characterization and validation

### Silicon and System Demo



System Validation Board



16 Gbps eye  
(socketed)

# Rambus Memory IP Solutions

Memory PHY Solutions for Networking, AI, Data Center and Automotive

## DDR4/3

- 3200Mbps
- x16 – x72-bits
- 1-4 Ranks
- DFI 4.0



## HBM2/E

- 2000 - 3200Mbps
- 1024-bit
- 2.5D design architecture



## GDDR6

- 12 - 16Gbps
- 2 x 16-bit channels
- Validated with Memory Controller


IN LAB

## DDR5 & HBM3

- DDR5: 4.8 – 6.4Gbps
- HBM3: 4Gbps

IN DEVELOPMENT/  
ROADMAP

Integrated tools for easy bring-up and characterization



LabStation Platform

- Easy-to-use PC Interface
- Interface to 3<sup>rd</sup> party software
- Pre-defined test scripts
- PHY control settings
- External instrument control
- System characteristics and analysis

*Rambus*

## SerDes PHY

David Kuo

Rambus Greater China

Sep. 12, 2019



# Speed, Low Latency and Reach

- HPC applications require a mixture of high-speed SerDes performance, low latency, electrical performance (reach in dB) and low power
- Depending on the application power/reach/latency/speed can be traded off to offer a best overall solution



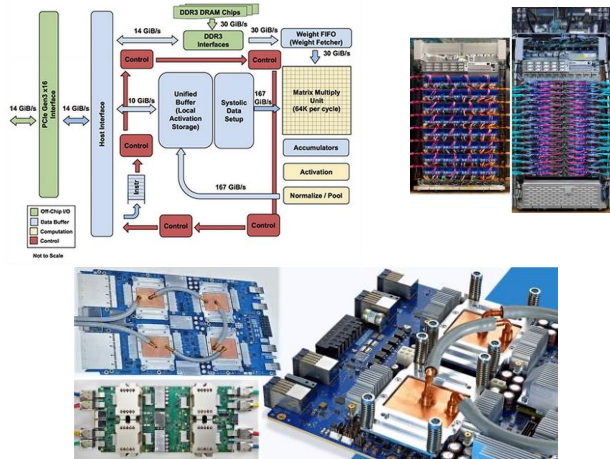
PCI  
EXPRESS®



# Applications – Contained Processing

ML, Mining and Chiplets


Google TPU Boards  
V1,V2,V3



Miners  
Avalon, Bitmain



Chiplets  
AMD Epyc Rome



Depending on the application we have a mixture of speed (e.g. Miners), low power (chiplets) and performance (rack-to-rack)

# Line Card Evolution

## SerDes Interfaces

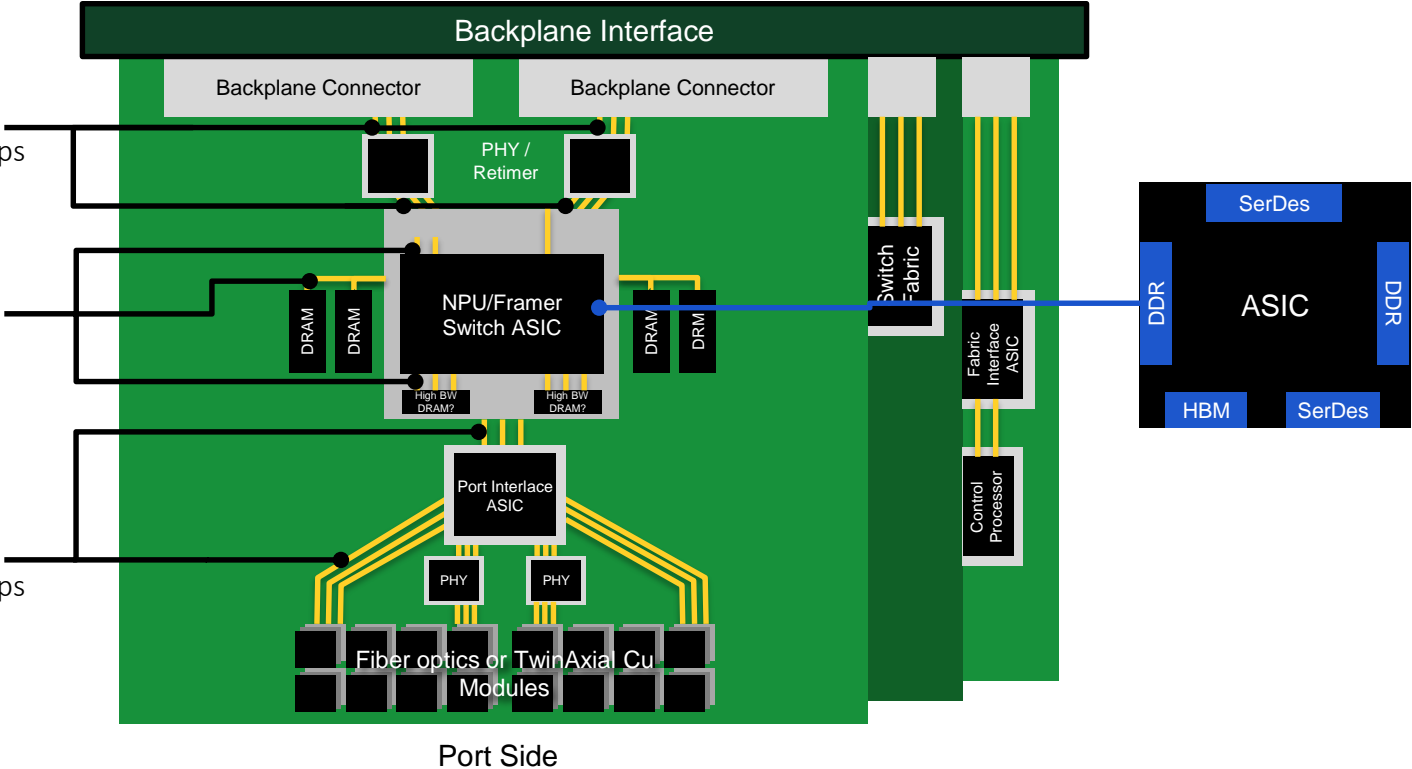
10/15/28Gbps → 56/112Gbps

## Memory Interfaces

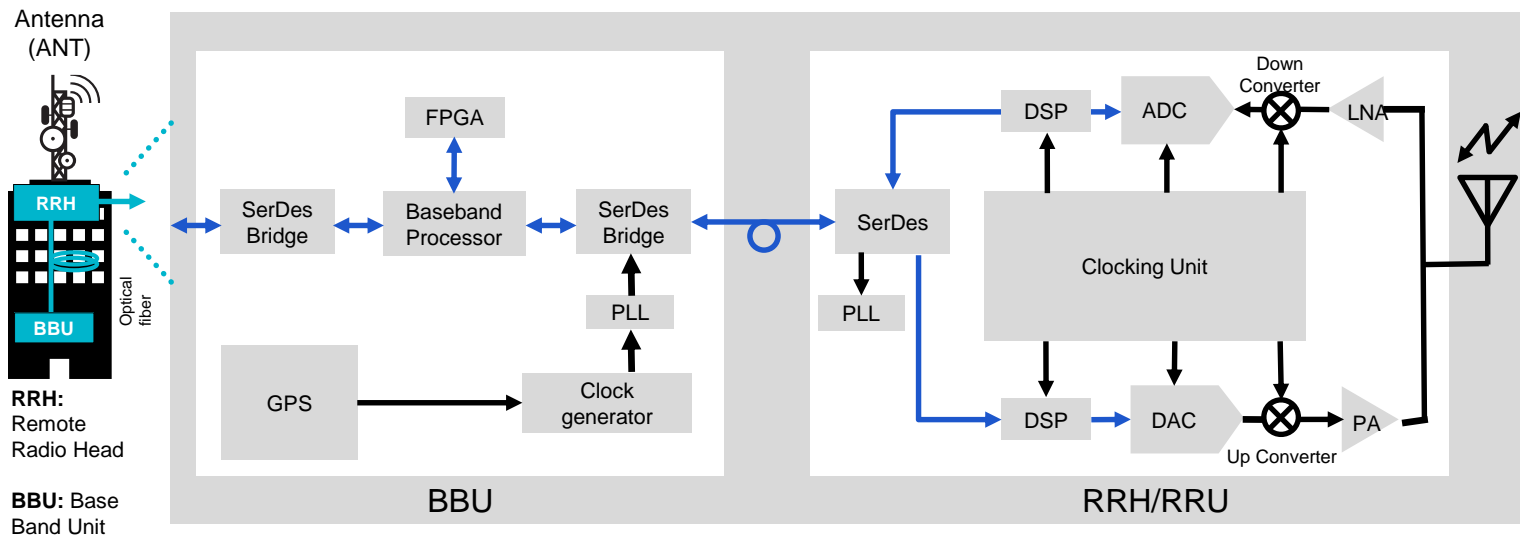
DDR4 → DDR5, GDDR6

## SerDes Interfaces

10/15/28Gbps → 56/112Gbps



# 5G: Wireless Infrastructure

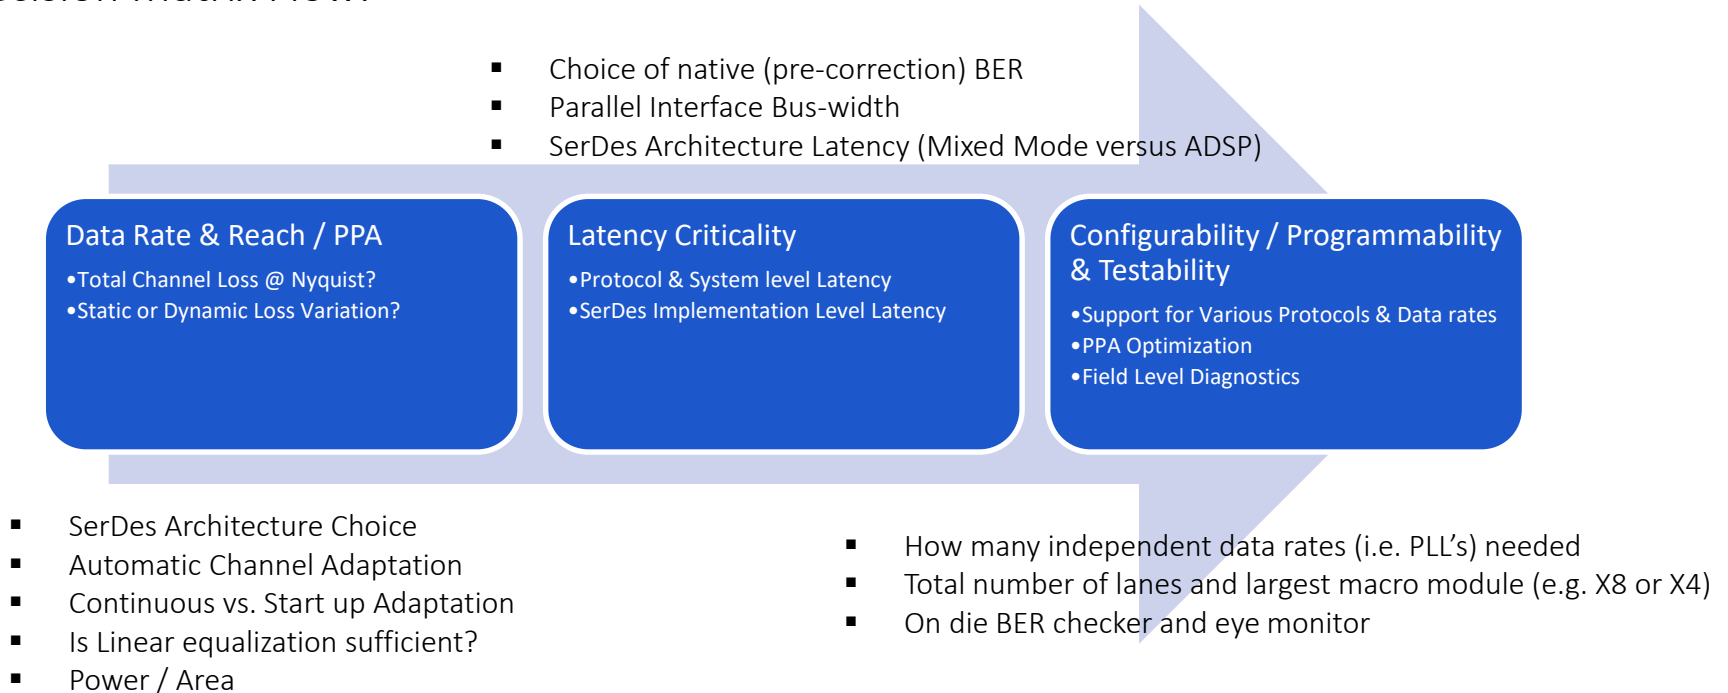


5G Deployment is the Key Driver in Wireless Space for next 3 years

- Higher link rate needs more local memory.
- CPRI from 12G → 24G → 48G → 56G
- JESD204 from 12.5G → 32G
- PCIe Gen3 → PCIe Gen4
- Ethernet from 25/28G → 50/56G
- DDR3 → DDR4/5 or HBM2

# SerDes Architecture for HPC Applications

## Decision Matrix Flow:



# Rambus SerDes IP Solutions

## Complete Solutions: SerDes PMA+PCS/MAC/Controller

### 1-28/32Gbps VSR/MR/LR

- PCIe 5\*/4
- OIF-CEI-28G
- Ethernet 25G
- JESD204B/C
- CPRI
- Interlaken



\*PCIe 5 in development

### 56Gbps LR

- 10-56Gbps
- Ethernet
- OIF
- CPRI



### 112/116Gbps XSR

- 28Gbps to **116Gbps** NRZ/PAM4
- TSMC 7nm Ultra high density SerDes
- Sub 1pJ/bit D2D

IN  
DEVELOPMENT

### 112Gbps LR

- 28Gbps-116Gbps NRZ/PAM4
- TSMC 7nm High performance SerDes
- 5pJ/bit

TAPED-OUT

## Integrated tools for easy bring-up and characterization





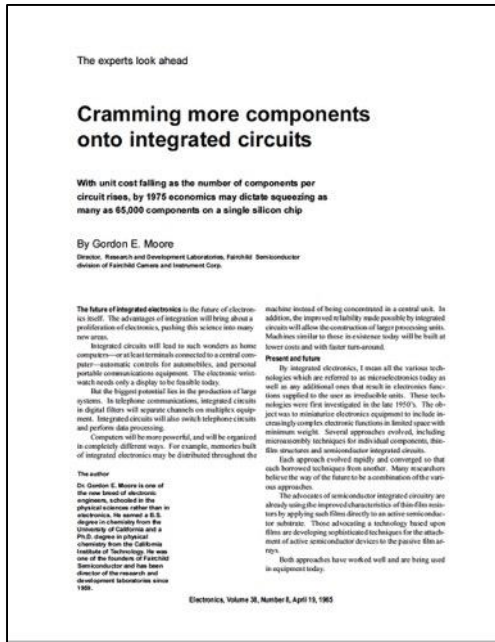
LabStation Platform

- Easy-to-use PC Interface
- Interface to 3<sup>rd</sup> party software
- Pre-defined test scripts
- PHY control settings
- External instrument control
- System characteristics and analysis

## Validated solutions with partners



# Everyone Is Familiar With Moore's Law from Gordon Moore's 1965 Paper





# On Page 4 of Gordon Moore's 1965 Paper...




“It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected. The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically”

AKA “System in Package”


# System in Package “Dis-Aggregation” Can Drive Performance and Control Costs





Today – Monolithic



Tomorrow – Modular



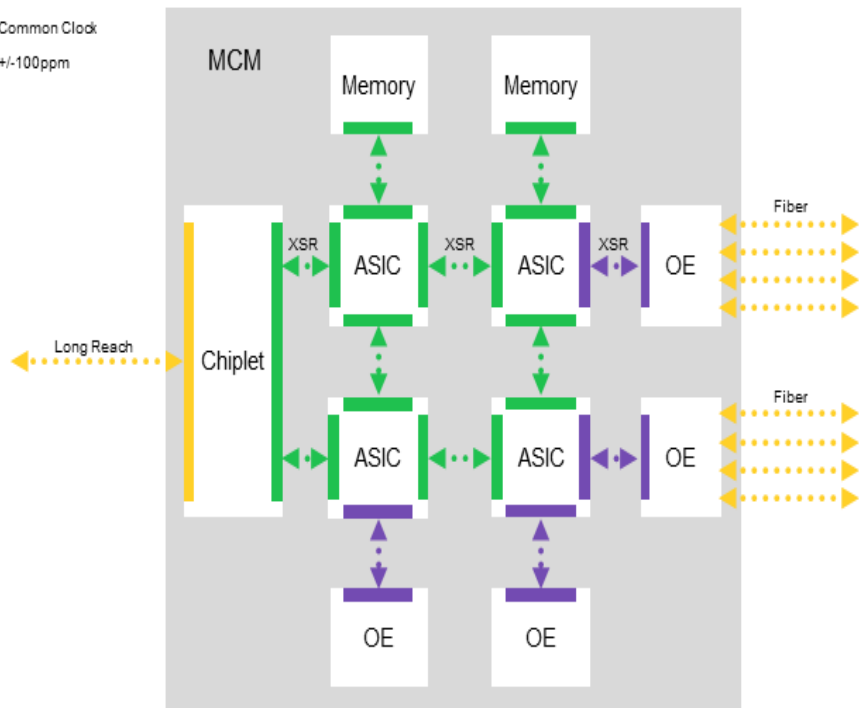
DARPA CHIPS




Open Domain-Specific Accelerator (ODSA) Workgroup

# Chiplets design


- ↔ XSR with Common Clock
- ↔ XSR with +/-100ppm



AMD Epyc Rome



Chiplets

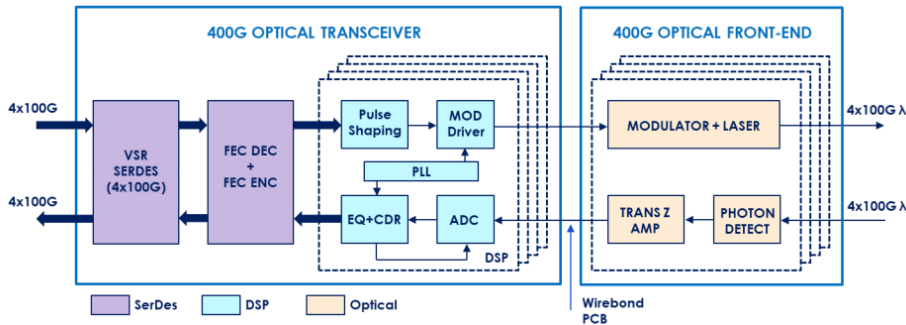


Beachfront: ~800s of Gbps/mm


# Evolution of the Optics Interface

- Shift away from convention optical modules due to power and cost
- COBO enables lower power due to shorter electrical link, but still presents a cost barrier
- Co-packaged Optics presents significant power savings due to very short electrical link and SiPho reduces component costs


## Convention Optical Module



## On-board Optics



## Co-packaged Optics



*Rambus*

Security

David Kuo


Rambus Greater China

Sep. 12, 2019



# The Growing Conflict Between Performance and Security

## 40 Years of Microprocessor Trend Data



Source: [www.karlsruh.net/wp-content/uploads/2015/06/40-years-processor-trend.png](http://www.karlsruh.net/wp-content/uploads/2015/06/40-years-processor-trend.png)

- Performance driven by Moore's Law
- More transistors = more performance and features, but also increases system complexity
- Security flaws often result from unexpected interactions
- As CPU complexity increases, security vulnerabilities grow exponentially

Attackers only need to find and exploit one vulnerability

# Hackers Used An IoT Connected Fish Tank Thermostat To Gain Access to a Casino's High-roller Database



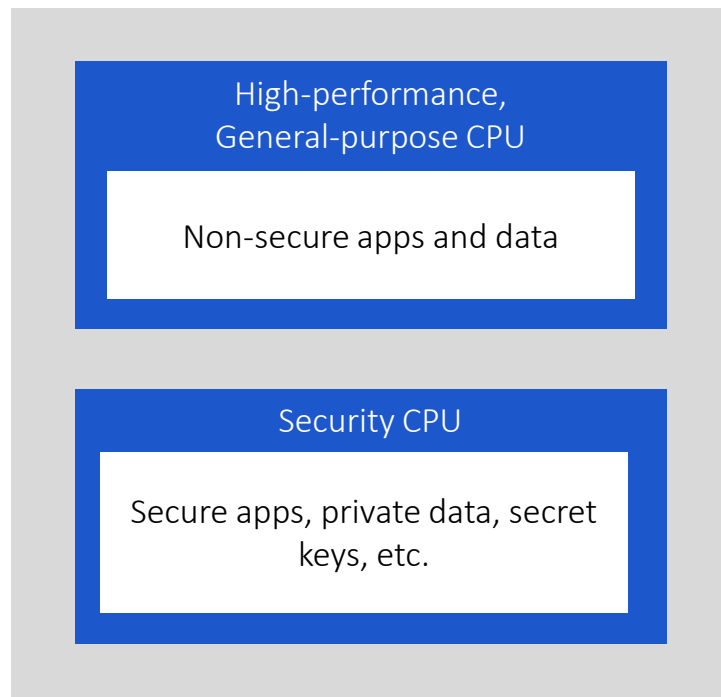
“The attackers used that to get a foothold in the network,” she explained. “They then found the high-roller database and then pulled that back across the network, out the thermostat, and up to the cloud.”



Darktrace CEO Nicole Eagan

# Improving Security with Siloed Execution

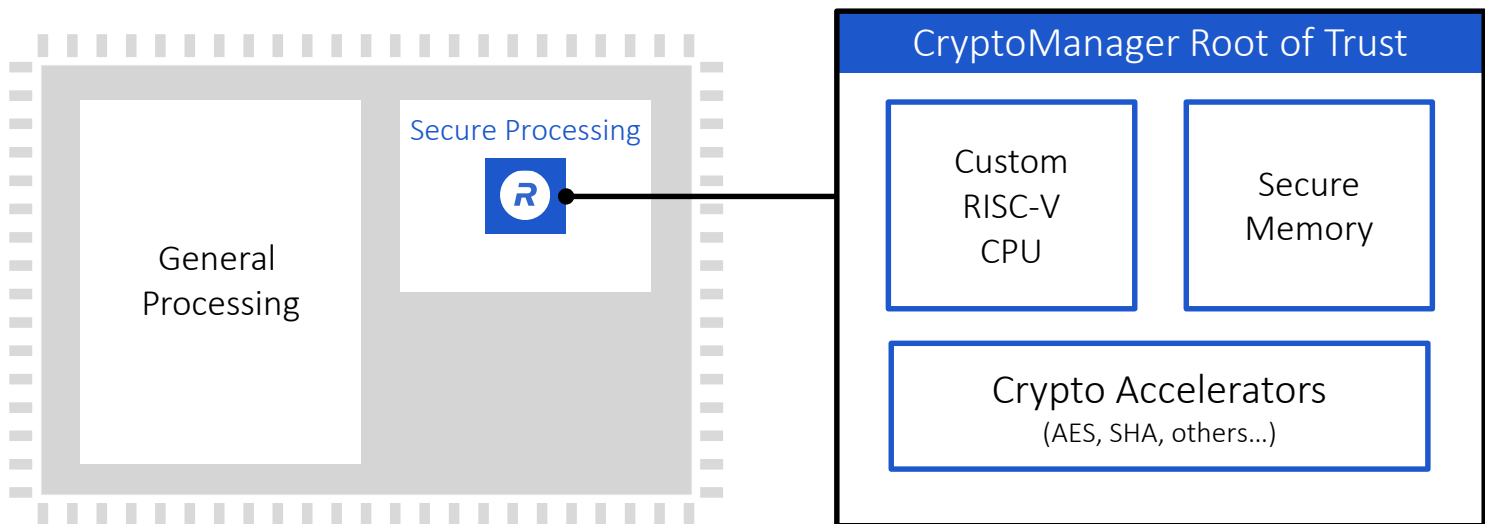
- New approaches needed to improve security in increasingly complex CPUs
- **Siloed execution:** Physically distinct CPUs that separate secure operations from operations that need to be fast
- **General-purpose CPU:** As fast, feature-rich, and complex as needed
- **Security CPU:** Keep the security CPU simple and optimize for security, stores keys and other security assets so they are not compromised even if the general purpose CPU is hacked



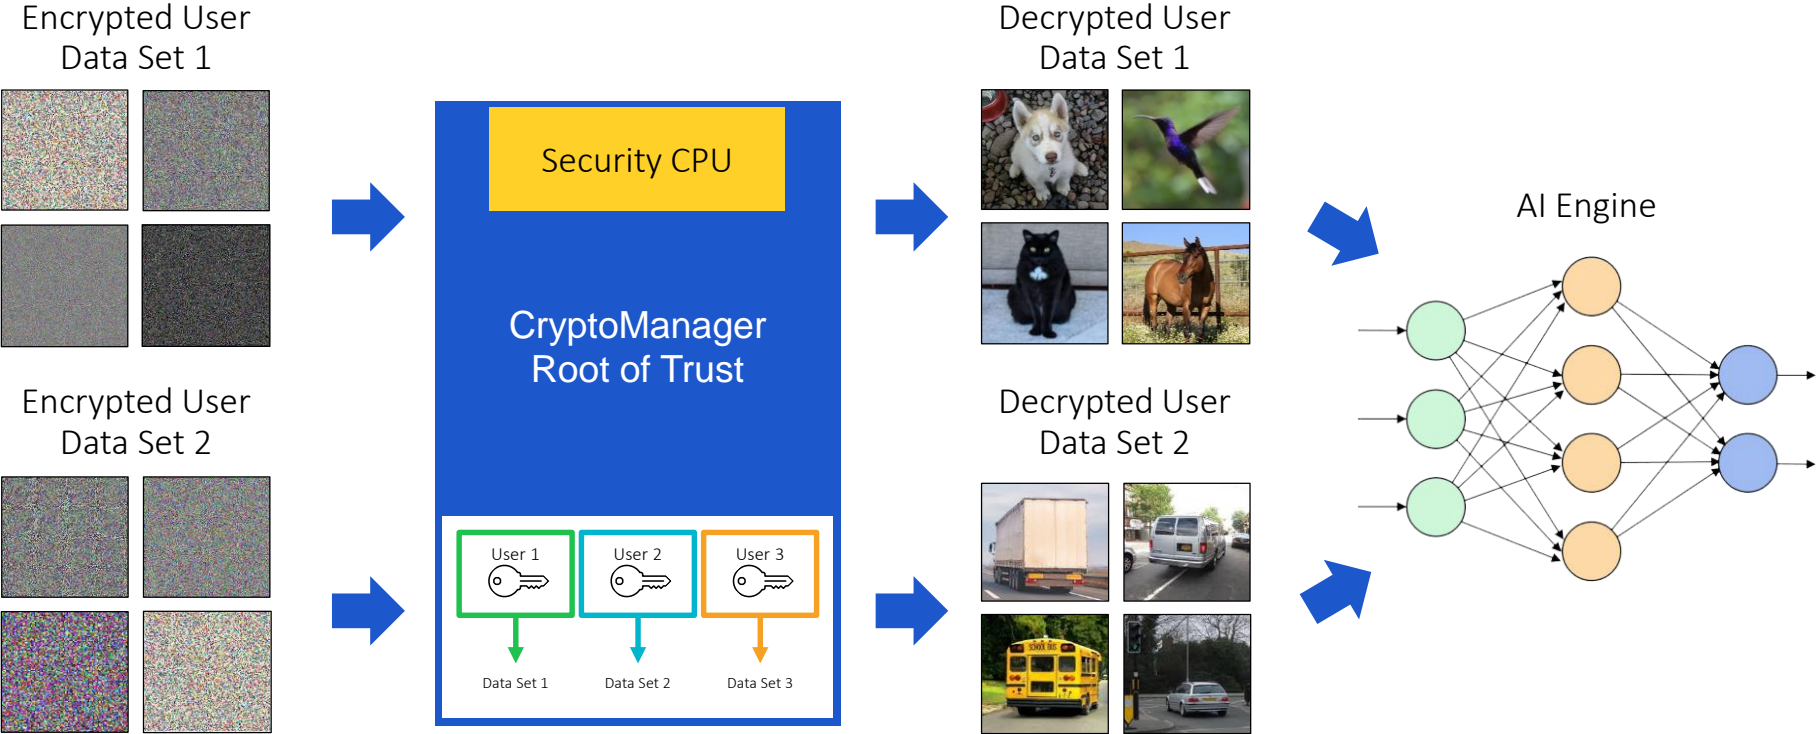
# CryptoManager Root of Trust

## Secure Functionality:

- Secure Boot
- Remote Attestation
- Authentication
- Runtime Integrity
- Key Vault




# Reducing AI Vulnerabilities with a Security CPU



CryptoManager Root of Trust Enables Multiple Roots of Trust in One Secure Processor

# John Hennessey, “From Now On, Must Treat Security as a First-class Design Goal”



“Security is more important than ever!”

From “The Era of Security”, John Hennessey (Chairman BoD Alphabet/Stanford University President Emeritus) in the HotChips 30 Plenary Keynote “Spectre/Meltdown & What it means for future design”



FORESHADOW



SPECTRE

# Summary

- Modern applications are changing the way we process information
  - 5G, growing digital data offer new opportunities for processing at the edge
- Memory and networks are critical bottlenecks to performance and power-efficiency
- Several challenges for building reliable high-speed interfaces
  - Accurate modeling for growing number of physical effects
  - Signal integrity and system engineering at critical for achieving high reliability
- Data and insights are increasingly valuable, security becoming more important
  - Must ensure processors and systems are secure by design
  - Need new approaches to address security as complexity grows



Thank you

**Rambus**  
Data • Faster • Safer