

MEXT 

Radical Reduction in Computing Costs


With AI-Powered Predictive Memory

First: Let's Look Back

Taking a walk down memory lane...


Moore's Law Left Memory Behind!

While Processing, Storage, and Networking Performance Have Exponentially Improved



The Evolution of the Processor

First Processor:
Single-Core, Simple In-Order Execution, CISC



1971

Superscalar, Mainstream Floating-Point



1993

64-bit x86 Architecture –
Basis for Modern Servers



2003

Multi-Core Processing for
Increased Efficiency




2006

AI Acceleration Enters CPUs,
Heterogenous Compute



2017

Chiplets and 3D
Packaging
Technology



Today

The Evolution of Networking

5G, Edge Networking,
Terabit Backbones

Broadband and Mobile
Networking Revolution, High
Speed Ethernet Becomes
Backbone for Cloud Storage




LAN on Motherboard



Birth of the
World Wide Web,
10BASE-T Ethernet



TCP / IP Becomes
Internet standard



ARPANET: First packet-
switched network



1969

1983

1991

2001

2010

Today

Revolutionary Shifts!

The Evolution of Storage

IBM 305 RAMAC:
First Hard Disk Drive



1956

CD-ROM Standardization
by Sony & Philips




1984

First Consumer
Flash Devices



1999

Enterprise Solid-State
Drives Arrive



2007

Cloud Storage
Becomes Mainstream



2010


QLC NAND, NVMe,
ZNS




Today

The Evolution of Memory


1 Transistor +
1 Capacitor Per Bit




1968




1990s



2000s



2010s

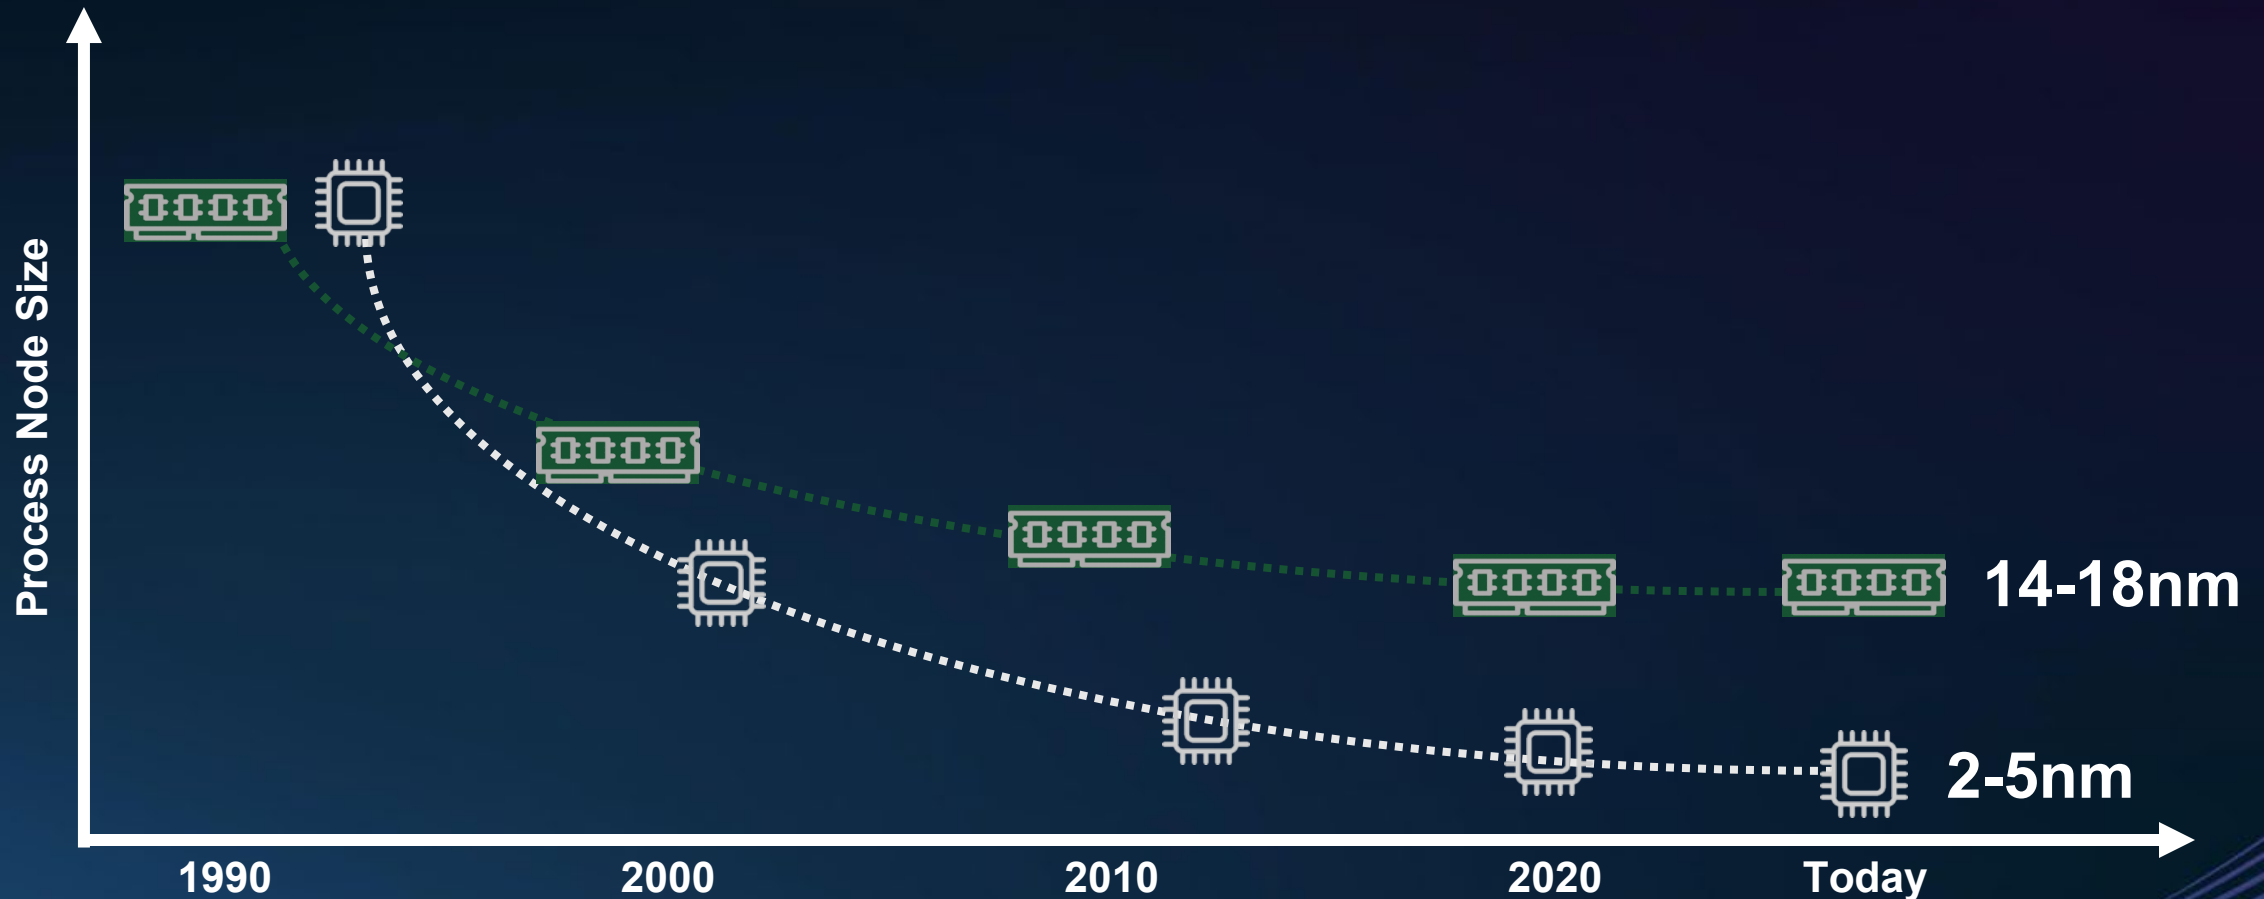


Today

Cell unchanged!

1 Transistor + 1 Capacitor Per Bit
*(w Higher Density, Faster Interfaces, and
Advanced Packaging)*

DRAM Process Technology Lags Behind Logic



“Software expands
to fill the available memory”

Wirth's Law

The **DRAMA** that **DRAM** Brings

Cost, capacity, and utilization problems

3 Major Memory Challenges

Cost per GB

DRAM is one of the costliest data center components

Capacity Scaling

Limited by the physics of storing reliable charge in ever-smaller capacitors

Utilization

Most memory pages in DRAM are cold most of the time

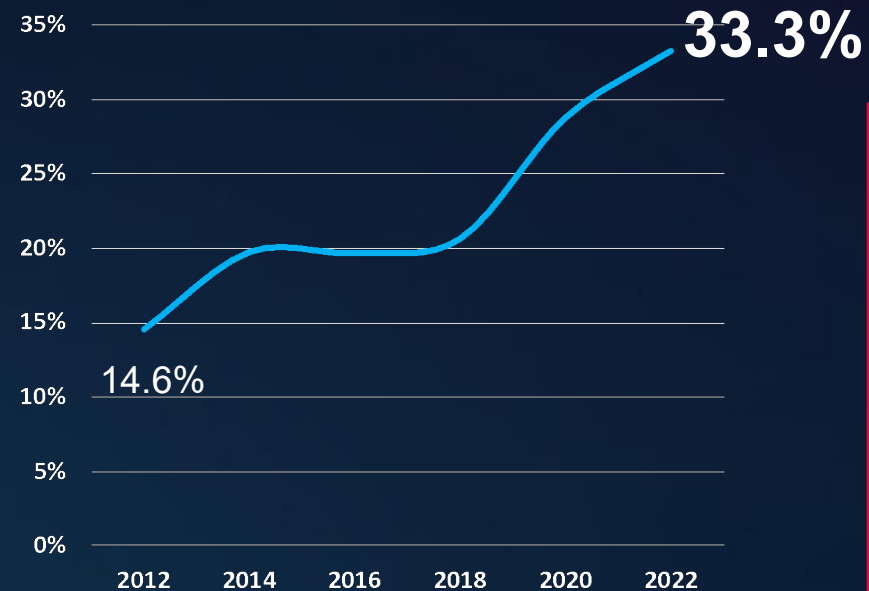
3 Major Memory Challenges

Cost per GB


DRAM is one of the costliest data center components

 Meta

DRAM is a **Huge Cost** in Meta's Data Center




DRAM is **50%** of the cost of individual server



3 Major Memory Challenges

Capacity Scaling

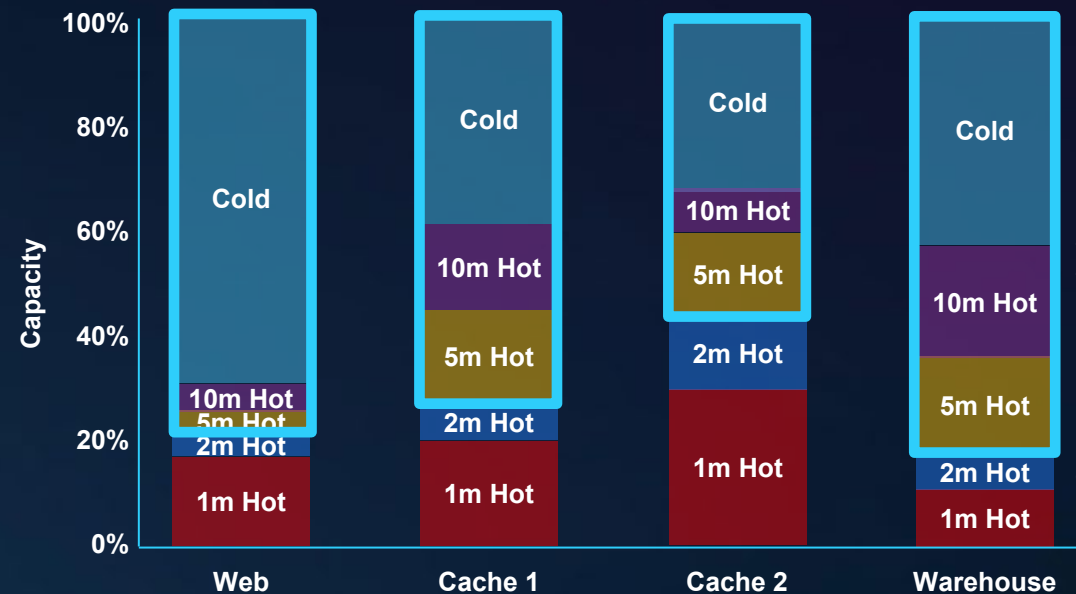
Limited by the physics of storing reliable charge in ever-smaller capacitors



3 Major Memory Challenges

Utilization

Most memory pages in DRAM are cold most of the time



Source: Meta paper on CXL tiering

Memory is often $< 50\%$ utilized

We've Seen Utilization Challenges Before

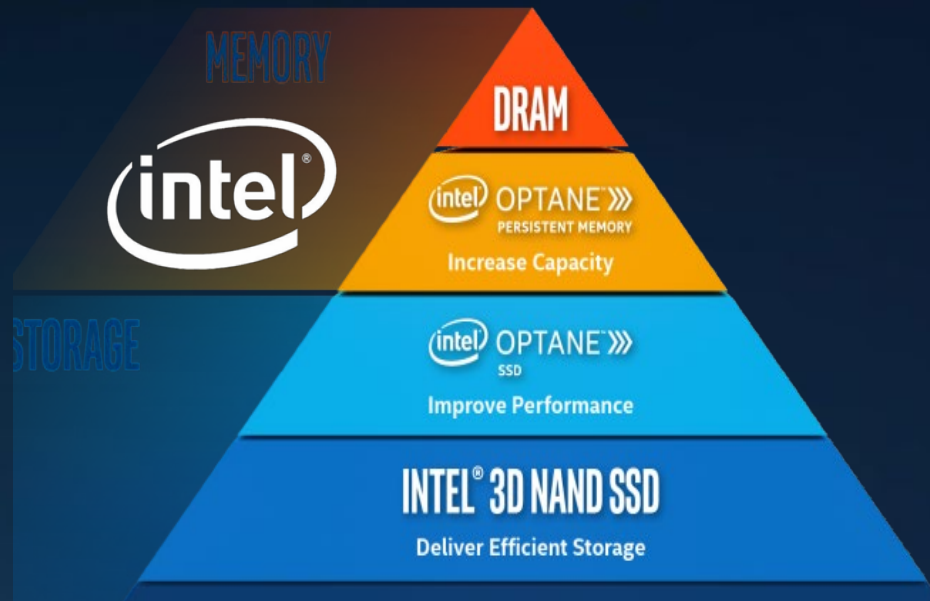
VMware solved the CPU
utilization problem...



Who will solve it for memory?



Previous Attempts to Solve Memory Challenges



CXL Compute Express Link

But They Have Missed the Mark!



The document is a white page with the Google logo at the top. The title is 'A Case Against CXL Memory Pooling'. The authors are Philip Levis (Google), Kun Lin (Google), and Amy Li (Google). The abstract discusses Compute Express Link (CXL) as a replacement for PCIe and mentions that CXL memory pools in datacenter and cloud networks consist of a large pool of memory for multiple machines share. The introduction discusses memory as a significant component of datacenter and cloud systems and mentions that recent papers report its fraction of a server's cost is 10% for AWS [14] and 50% for Azure [21]. A large, bold, red 'NOT READY' stamp is written diagonally across the document.

How Do You Bring Down the Cost of Memory?



Increase DRAM Utilization

Eliminate overprovisioning and stranded DRAM



No Hardware or Software Changes

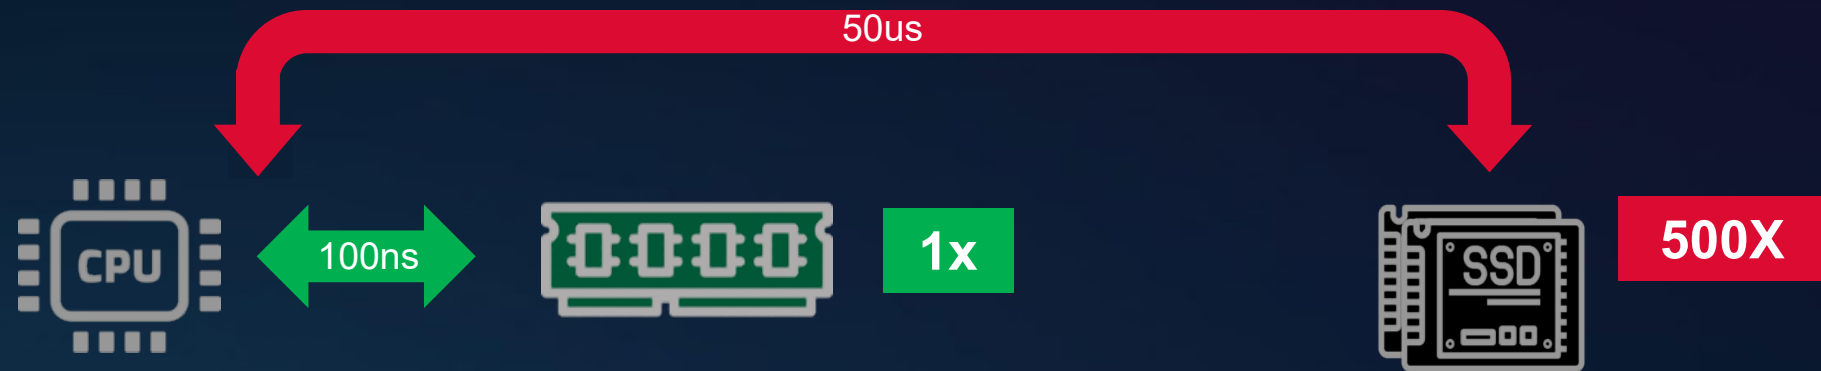
HW / architectural and SW changes drive up cost



Bring Flash into Memory Tier

Flash is 20X - 50X cheaper than DRAM & 30X lower power

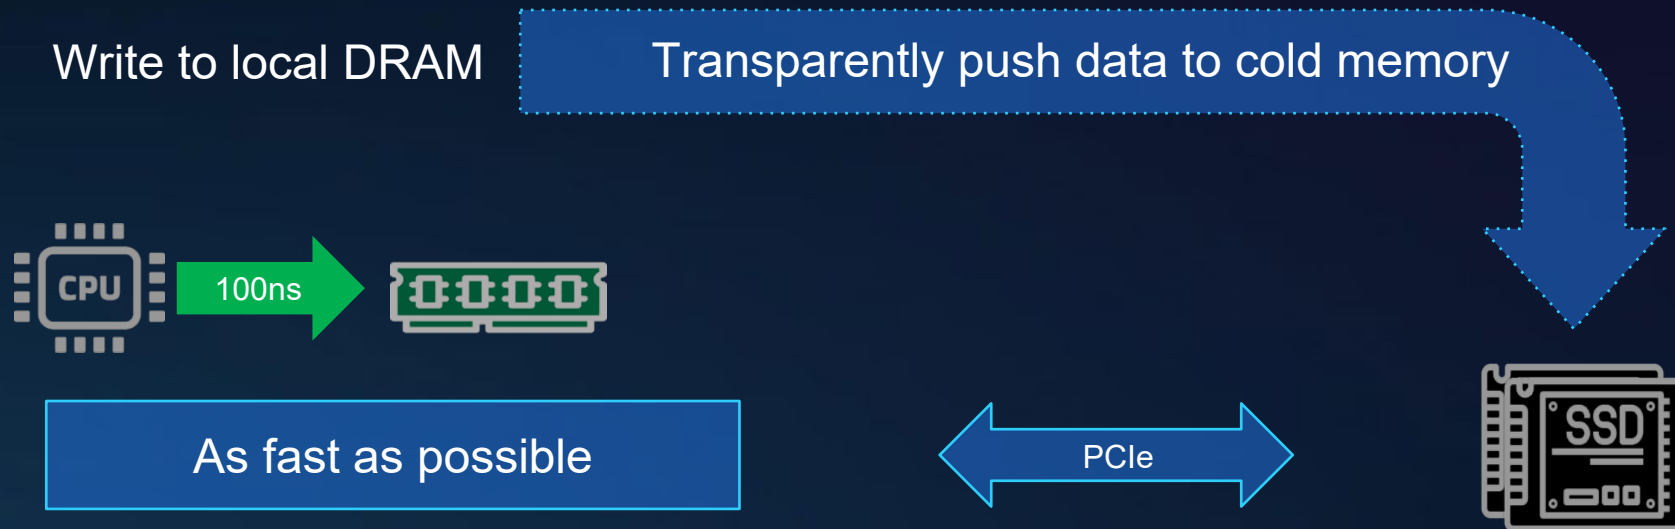
“Great Idea, But It Won’t Work...”



Meet MEXT Predictive Memory™

Radically reducing computing costs and expanding memory capacity

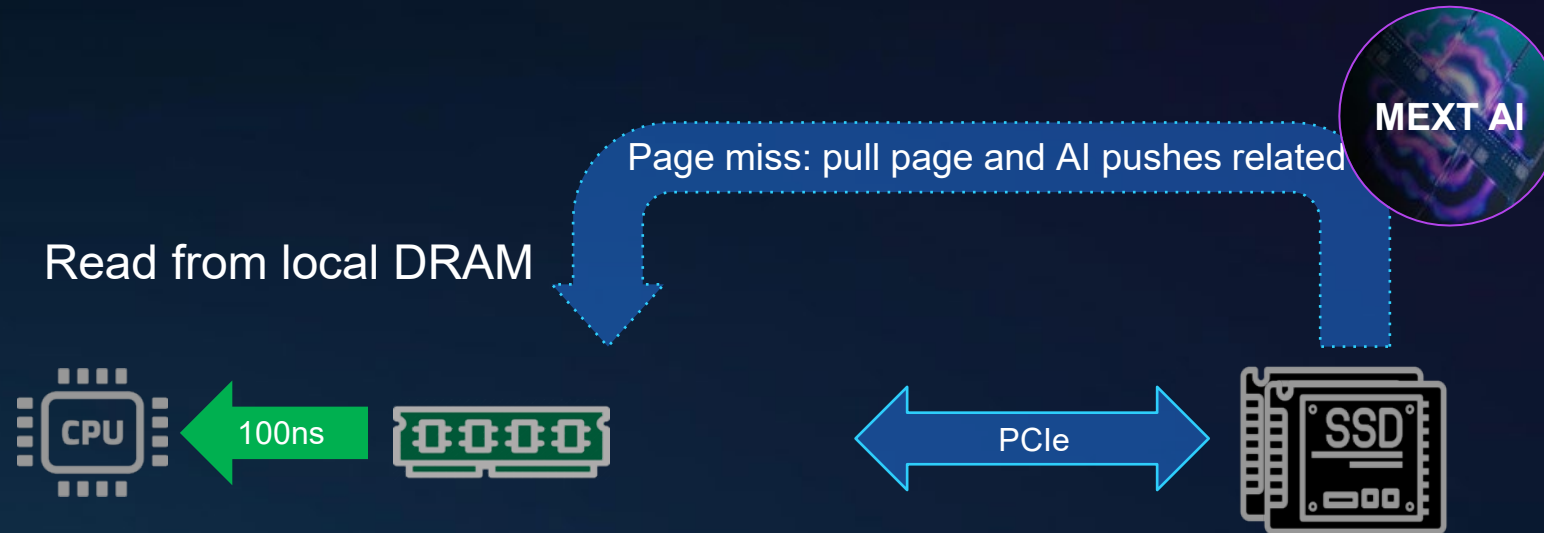
Writing to Memory Is Easy





MEXT Innovation: Predictive Memory

AI engine predicts and transparently tiers memory



- AI Push Engine predicts what will be needed in advance
- Places memory pages optimally before they are requested
- Goal: all but initial miss happen as quickly as possible

MEXT AI-Powered Predictive Memory

No HW / SW changes, use FLASH as memory



MEXT Plug-In

Lightweight, low-overhead Linux Driver

MEXT Predictive Memory Engine

Ensemble of AI models enables transparent auto-tiering of memory pages

MEXT View™ Observability Platform

Illustrates memory utilization and MEXT AI prediction accuracy

MEXT AI Engine Development


Heuristic

Savant
Leopard




LSTM

L1
L1A (with Attention)



Transformer

Transcend
T2 - Hybrid




Ensemble of AI models – continuous learning and self-optimizing

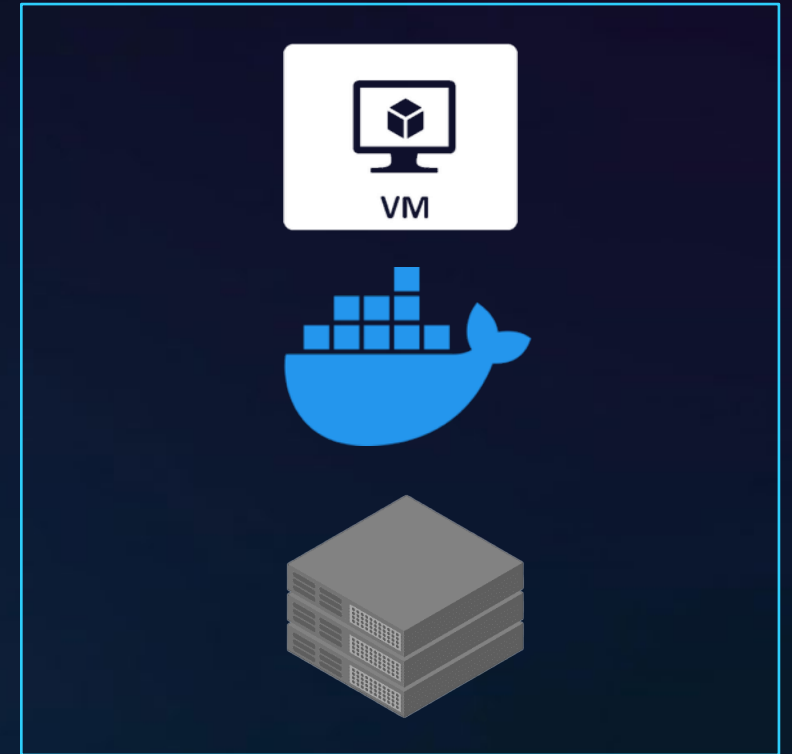
MEXT Software Deploys Anywhere in Minutes



**On-Premises or
Cloud Environments**




**Intel, AMD,
or ARM Processors**



**VMs, Containers, or
Bare Metal Configurations**



MEXT View: Visualization of DRAM Utilization



Performance Analysis

Delivering superior price-performance across various workloads

MEXT Workload Examples

Relational In-Memory Database



Analytics



In-Memory Key Value



Caching Key Value



Graph Database



AI-Native Vector Database



AI Large Language Model



Animation Rendering




Performance

Dollar

Performance Summary



Comparing 32GB DRAM & 32GB MEXT+Flash vs. 64GB all-DRAM System





AWS Cost Savings with MEXT

Memory	AWS – Today	AWS + MEXT 50% DRAM / 50% Flash	AWS + MEXT* 25% DRAM / 75% Flash	Customer Savings
256GB	\$1740 / month	\$893 / month	\$680 / month	49% - 61%
512GB	\$3490 / month	\$2207 / month	\$1149 / month	37% - 67%
1TB	\$8145/ month	\$4413 / month	\$2719 / month	46% - 67%
4TB	\$19,474 / month	\$11,785 / month	\$10,137 / month	39% - 48%

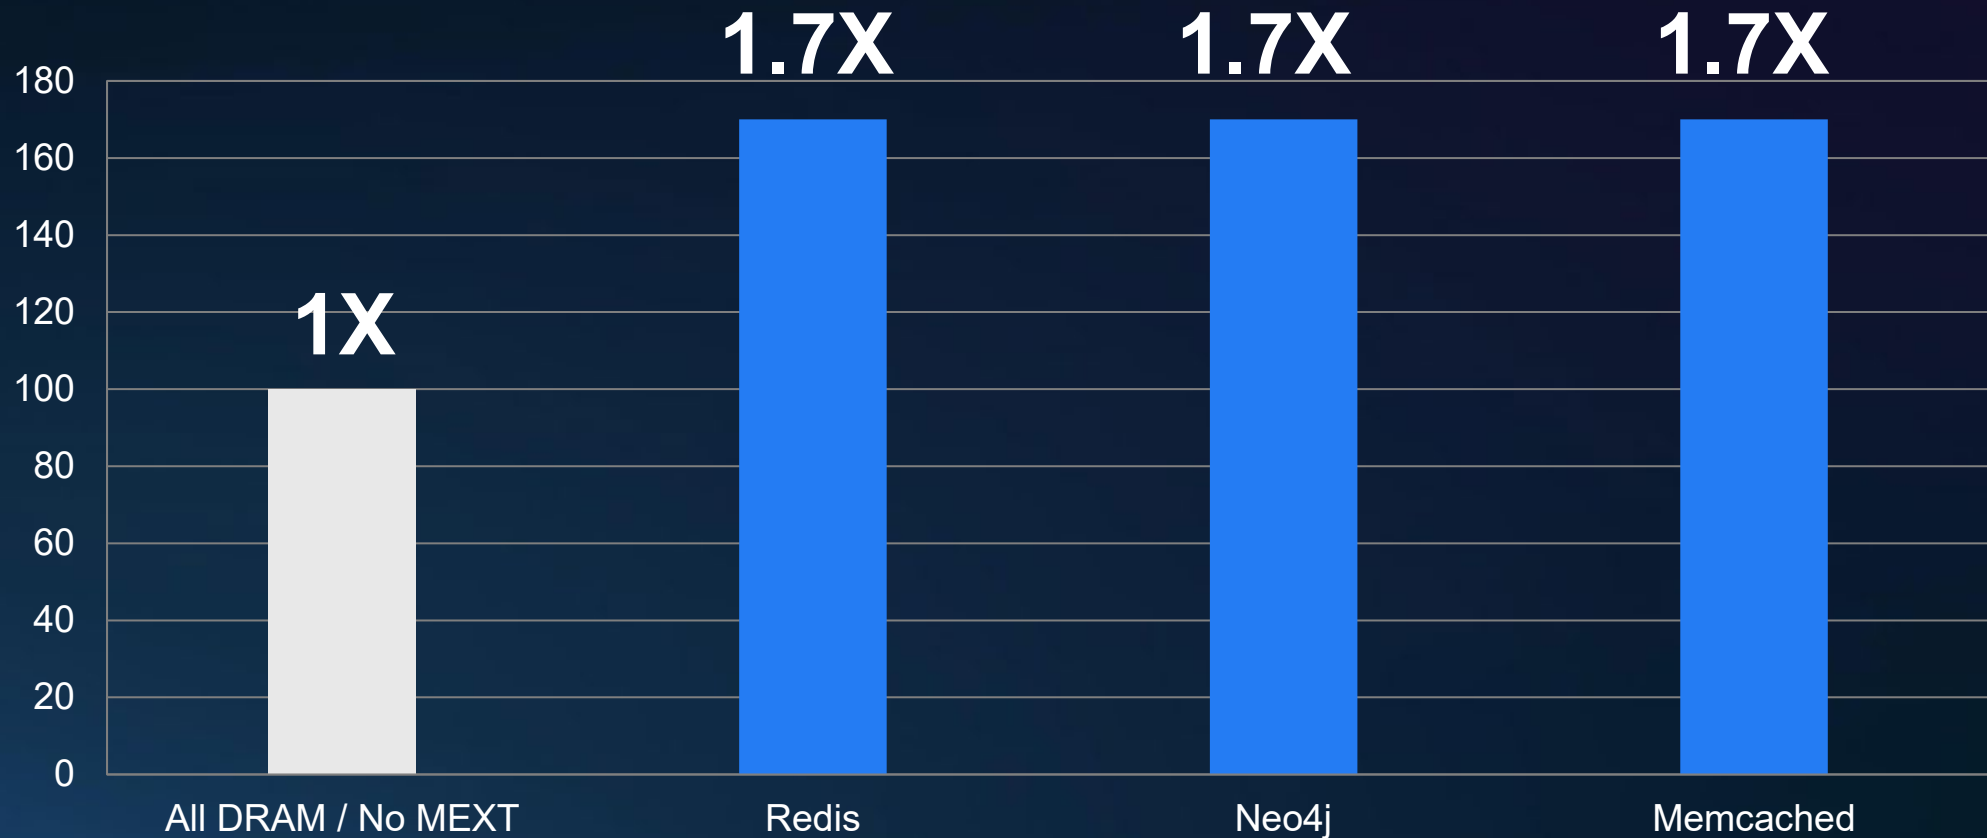
Includes MEXT

**37% - 67%
Savings**

Deploy in minutes. Save instantly.

Driving Superior Performance / \$


Comparing 32GB DRAM & 32GB MEXT+Flash vs. 64GB all-DRAM



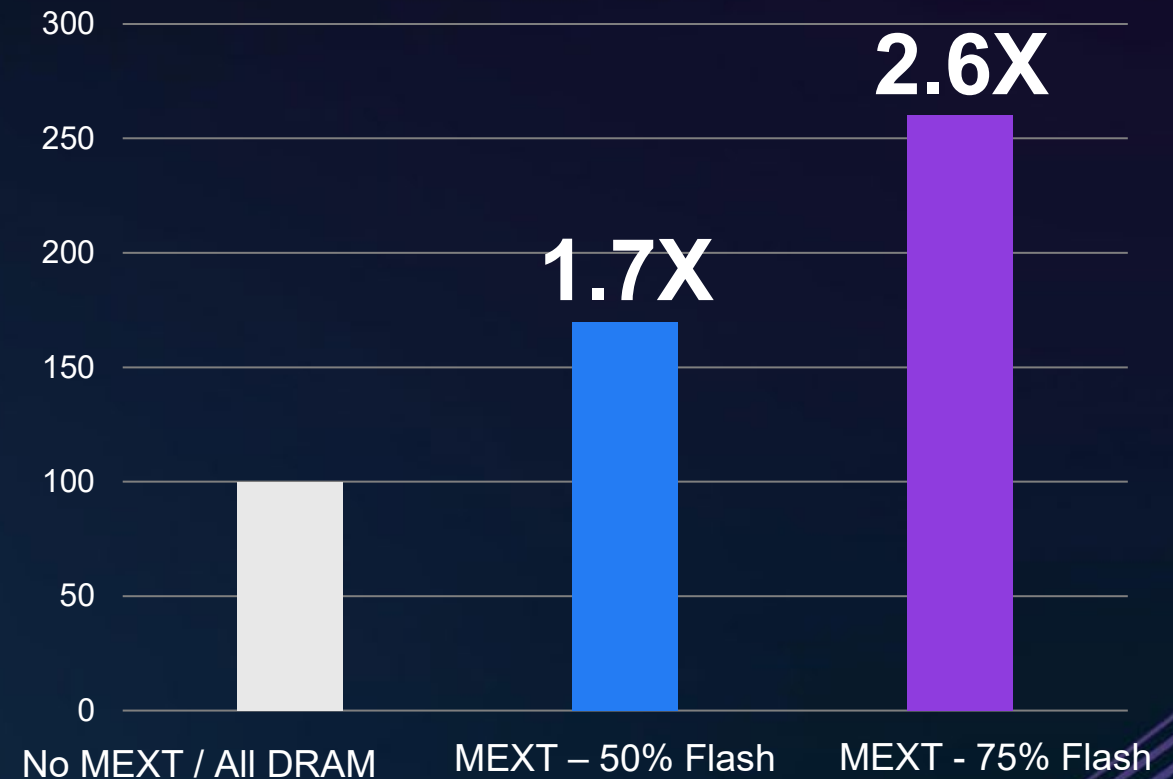
Improved Economics for Neo4j




Neo4j (WDBench) – 64 GB Memory



Neo4j Performance / \$





Power Reduction Through Server Consolidation



1000 servers

1 TB / server
(1 TB DRAM)



500 servers

2 TB / server
(1 TB DRAM & 1+ TB MEXT+Flash)

17,540 MWh
Power Reduction

6,924 Acres
Carbon Offset

\$15M
Cost Savings

Wrapping Up

The right solution, the right team, the right time

3 Major Memory Challenges: SOLVED

By AI-Powered Predictive Memory

Cost per GB

Leverages Flash as memory,
20-50X cheaper cost-per-bit

Capacity Scaling

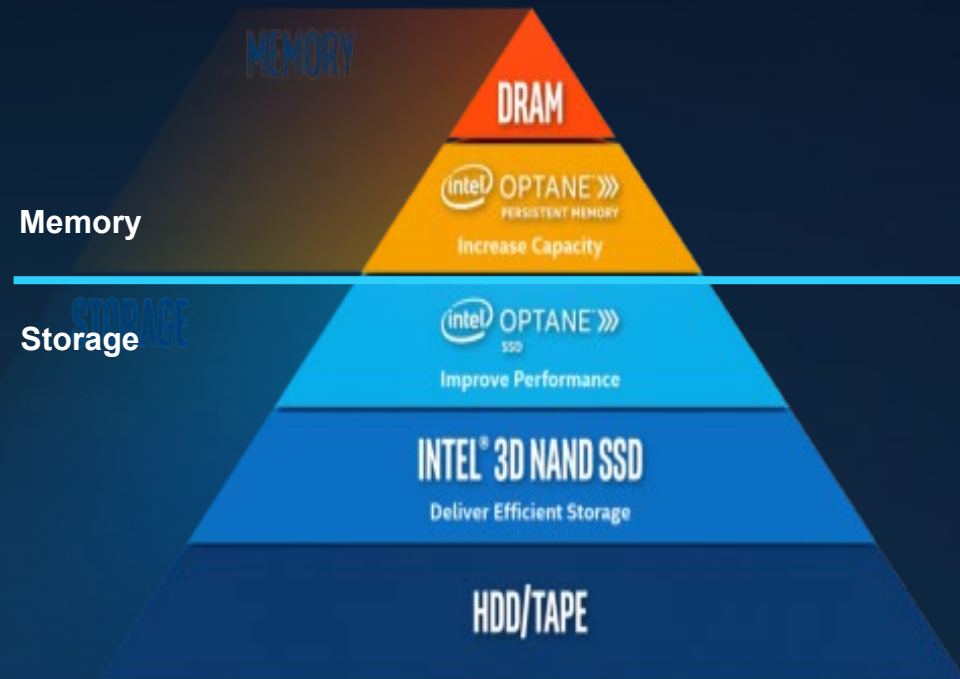
Flash capacity is virtually
boundless

Utilization

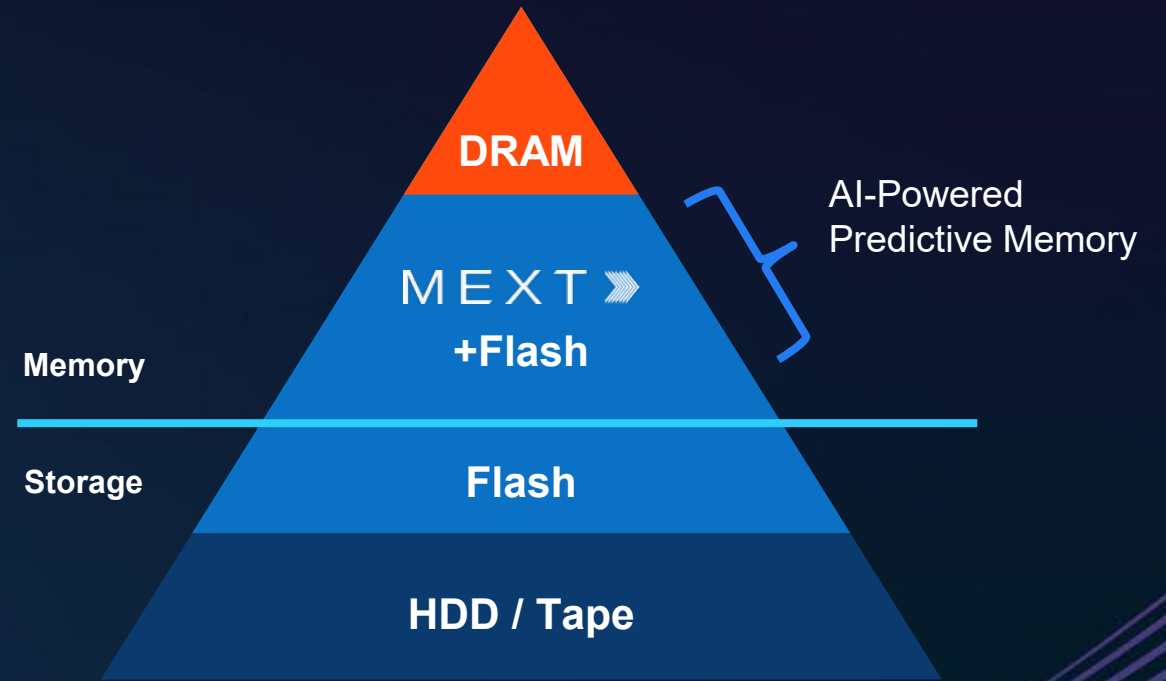
Intelligent transparent tiering
across DRAM and Flash

MEXT Brings Flash into the Memory Tier

Intel Memory Pyramid



MEXT Memory Pyramid





Industry Innovators and Pioneers

Gary Smerdon
CEO / Founder

Former CEO of TidalScale, exec at Fusion-io, LSI, AMD, industry visionary

David Reed
CTO / Founder

Chief Scientist at TidalScale & Lotus, SVP SAP, HP Fellow, recognized innovator, MIT PHD

Carl Waldspurger
Chief Scientist

Early @ VMware, led core innovations for 10+ years, 138 patents / 97 pending / 35 papers / 6 best-paper awards, MIT PHD

Michael Berman
Chief Engineer

Former Distinguished Technologist HPE, TidalScale Chief Engineer, CTO at Catbird Networks

Bobby Minnear
VP of Engr / Founder

Former TidalScale VP of Engineering, numerous Engineering lead roles

Adam Bane
VP of Product

Former TidalScale VP of Product, former Director R&D IBM Cloud

Greg Shomstein
VP Sales

Former Executive / Sales VP at PicsArt, TidalScale, IBM

Eric Fitzpatrick
VP Solutions Engineering

Former TidalScale VP of Solutions Engineering, former HPE, P&G, SAP

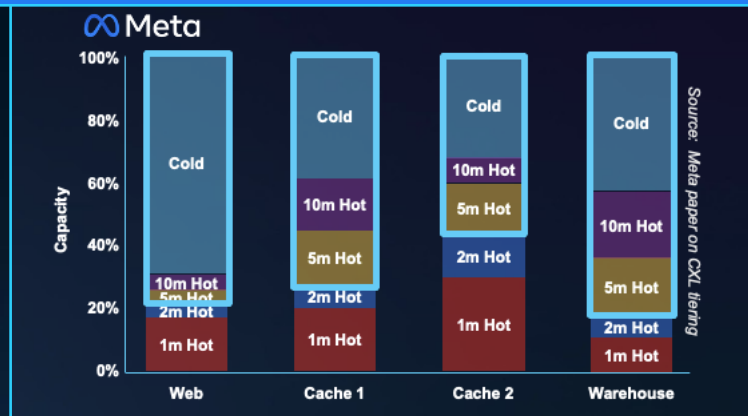
Advisors:

Fred Weber – Former CTO AMD, **John Williams** – GCP, Azure, AMD, **Steve Yatko** – FinTech Leader, **Toby Bloom** – Bioinformatics Expert

MEXT Partners



The Pain Is Real



The Urgency Is Growing

AI ↗

MEXT Is the Solve

Without MEXT



MEXT →

With MEXT



M E E X T

