

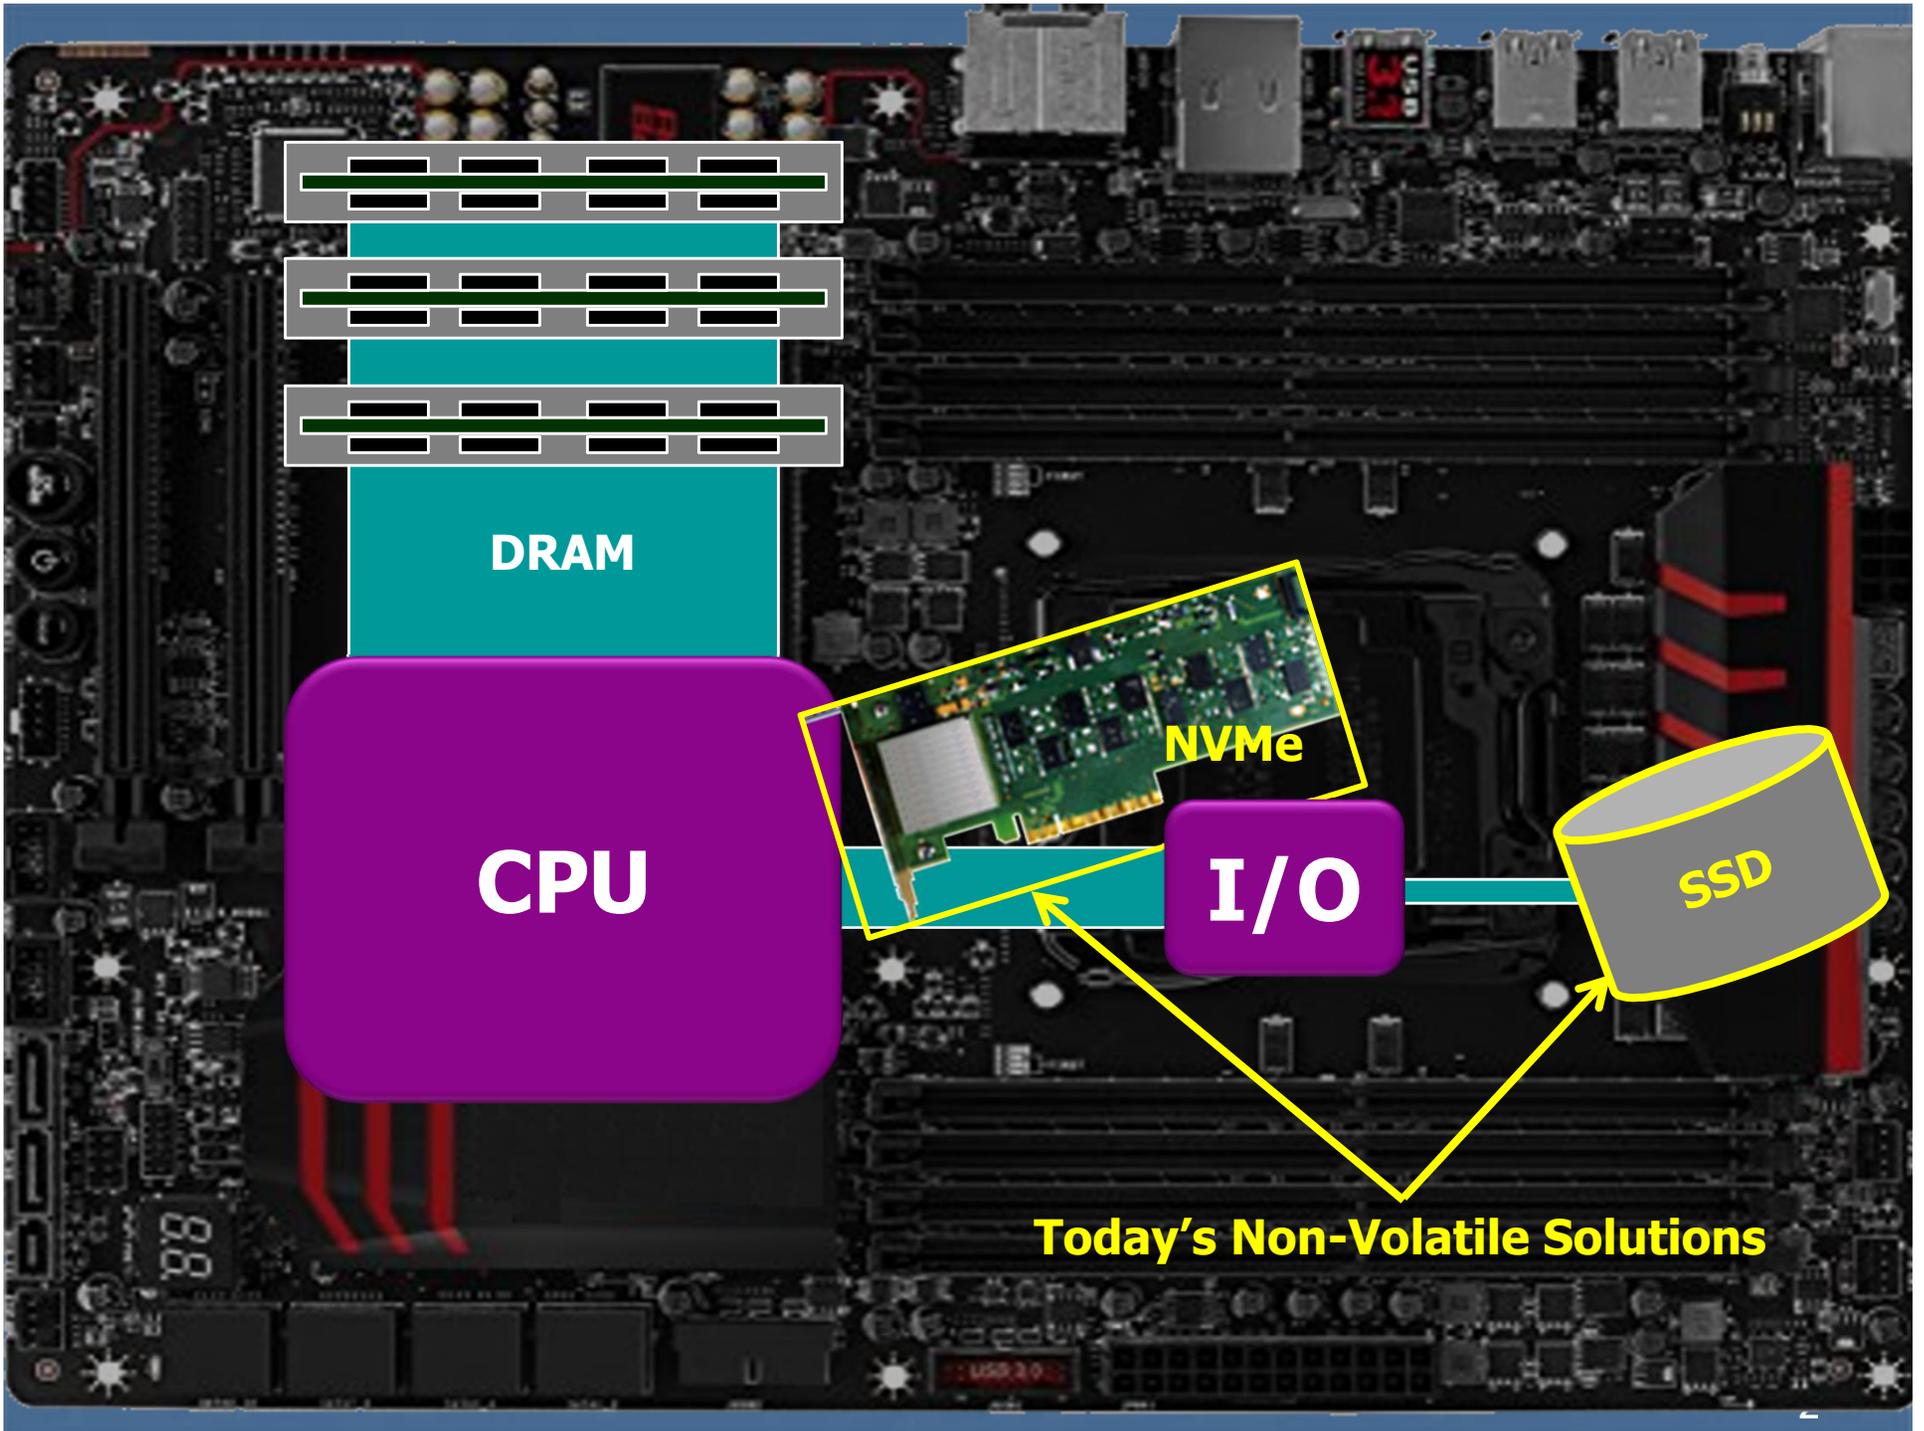
The JEDEC logo is displayed in a bold, dark green, sans-serif font. A red horizontal line is positioned below the letters 'E' and 'D', extending to the right edge of the logo's bounding box.

# Towards an Heterogeneous Memory Channel with Hybrid Modules

Bill Gervasi  
October 2015



**DISCOBOLUS**  
DESIGNS



**DRAM**

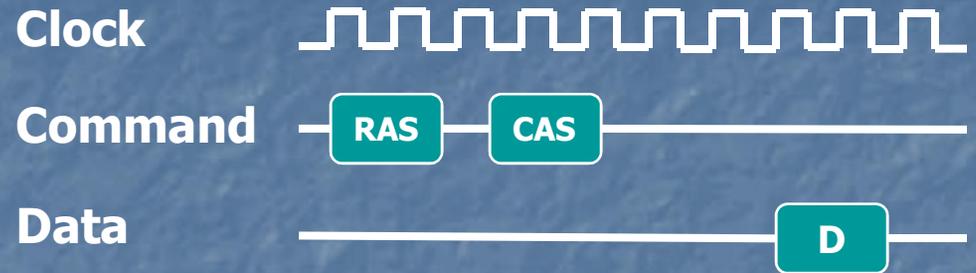
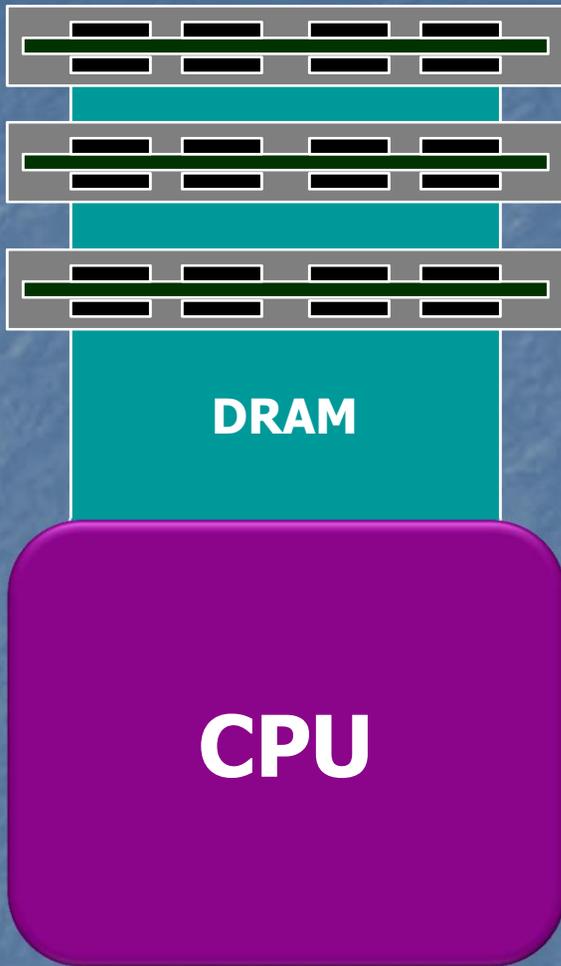
**CPU**

**NVMe**

**I/O**

**SSD**

**Today's Non-Volatile Solutions**



# DRAM Interface

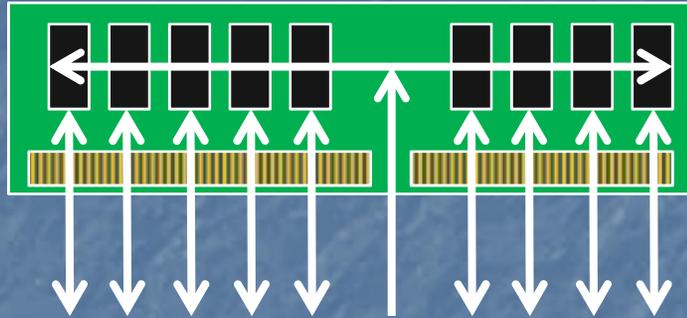
"RAS-CAS" command protocol

Deterministic data

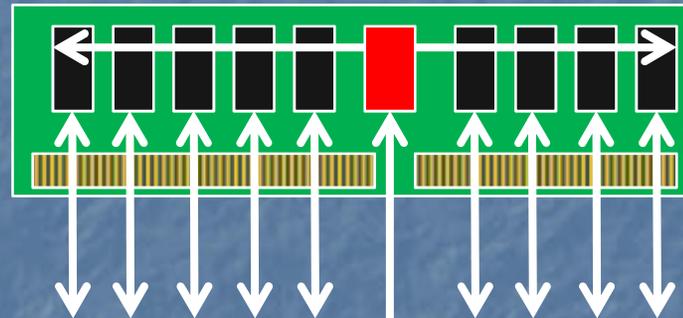


# Evolution of DRAM Interfaces

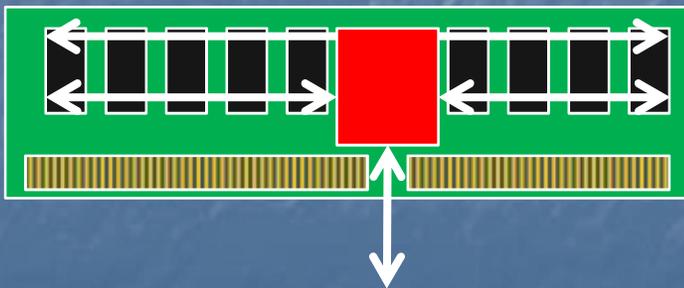
Unbuffered



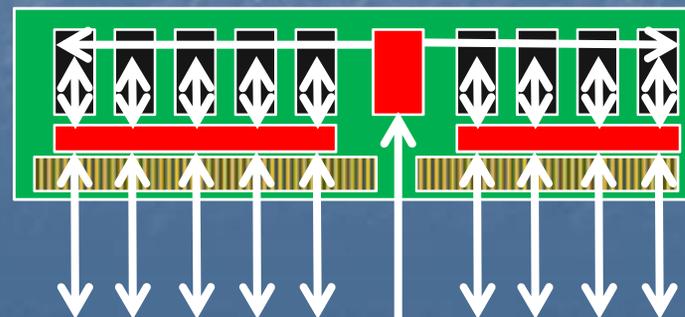
Registered



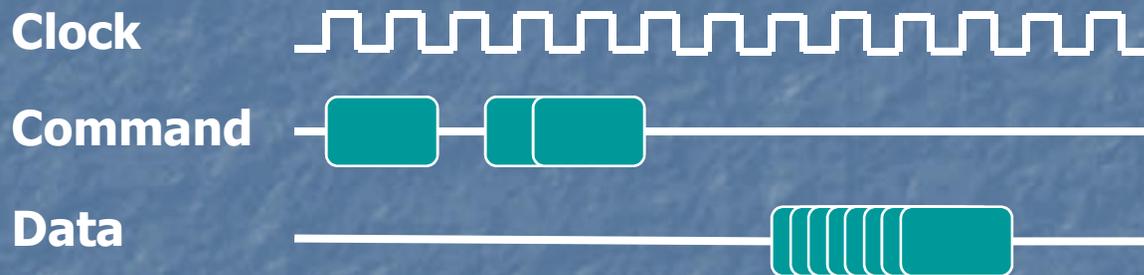
Fully Buffered



Load Reduced



## Each evolution maintained RAS-CAS



Minor “tweaks” to accommodate differences

e.g., Clock delay due to register

...but... no longer just a pure DRAM interface

**RAS-CAS is just a “protocol”**



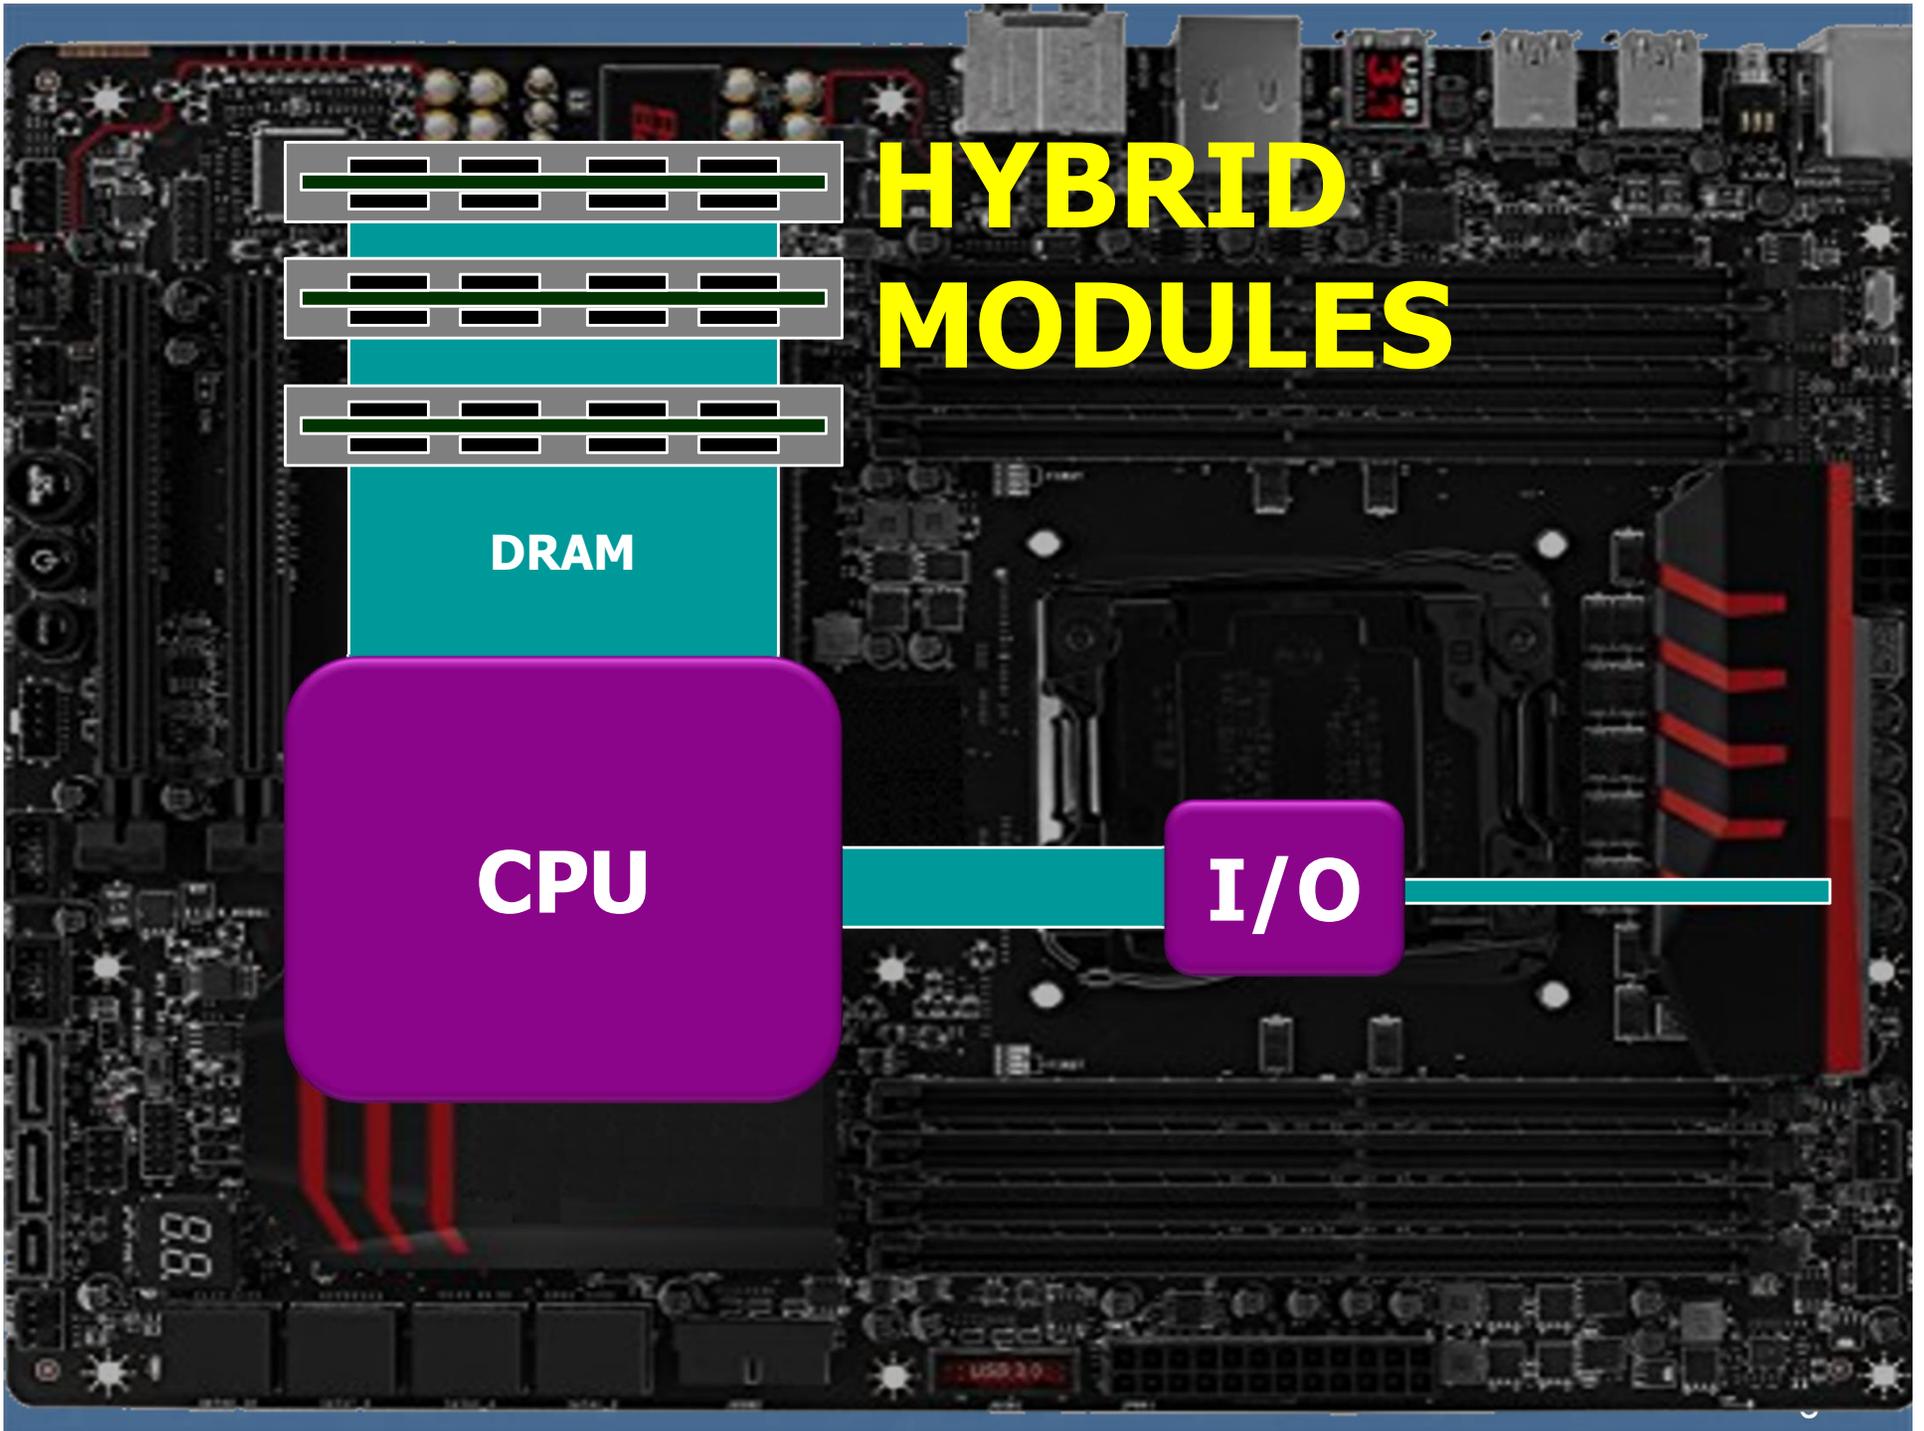
# HYBRID MODULES



DRAM

CPU

I/O



## Legacy Storage

## Leading Edge Storage

Hot Data

Latency

Tier 0

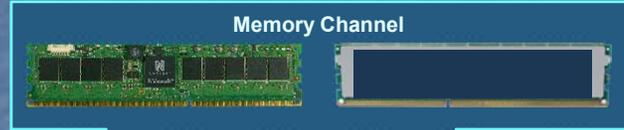


Access to Host

Caching  
Indexing  
In-memory

Hybrid Modules

Memory Channel



.01  $\mu$ s

Tier 1



Analytics OLTP  
Database Web Hosting

NVMe

PCIe



10  $\mu$ s

Tier 2



Network Accessible

Mail Servers CRM  
VOD Media Streaming Surveillance



100  $\mu$ s

Tier 3



Data Warehouse Content Delivery  
Archive Backup

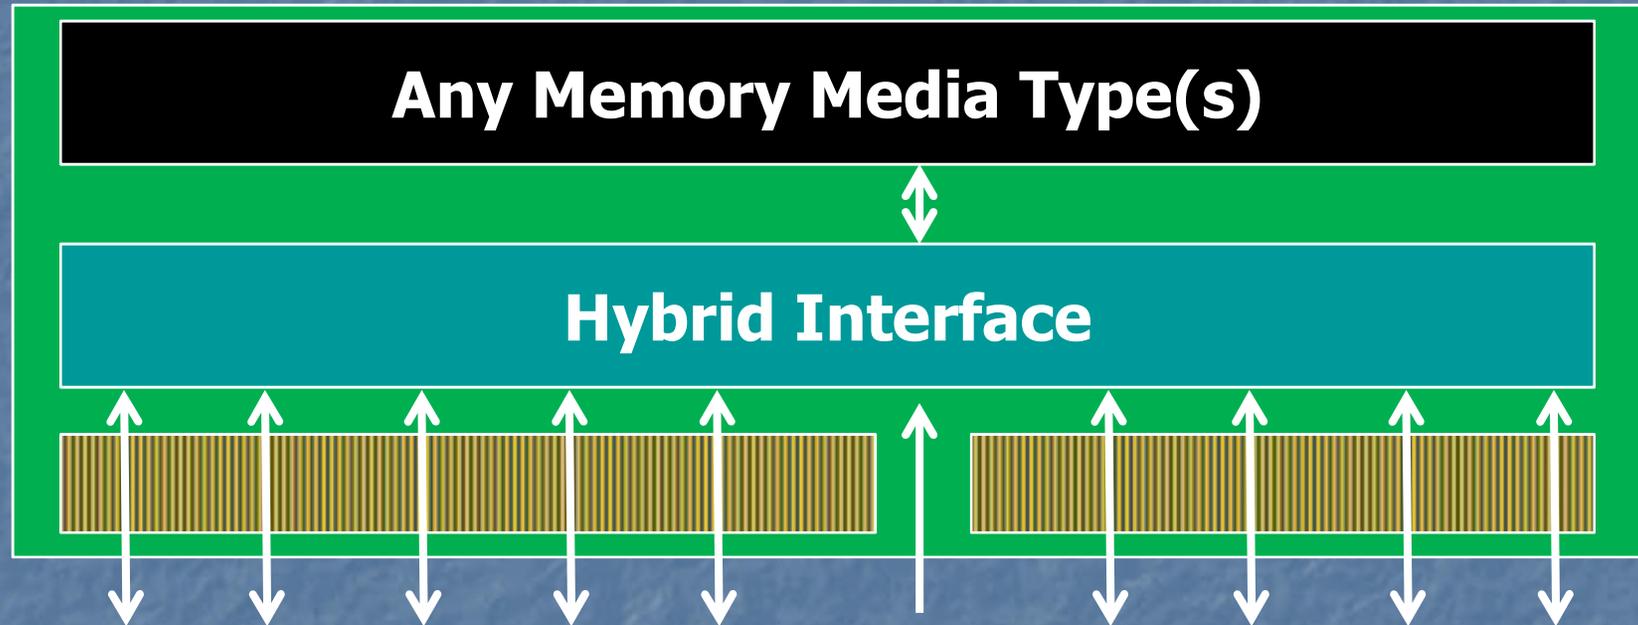


1,000  $\mu$ s

Cold Data



# Hybrid Modules

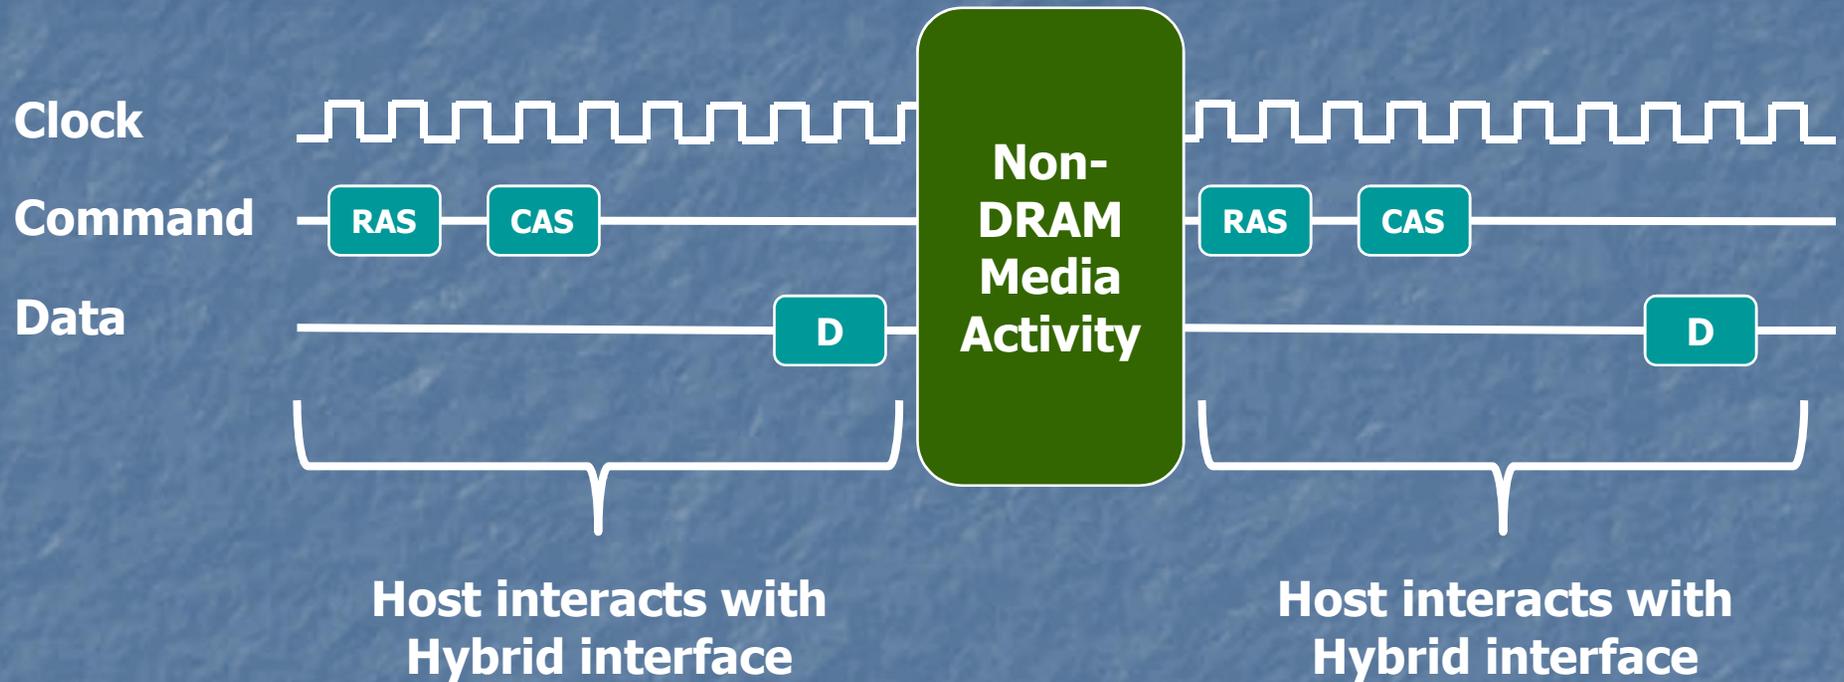


**Presents a DRAM-compatible RAS-CAS interface**

**One or more non-DRAM memories behind**



# Hybrid Interface



Still the DRAM compatible "RAS-CAS" command protocol



# **NVDIMM: Hybrid with non-volatile media**

---

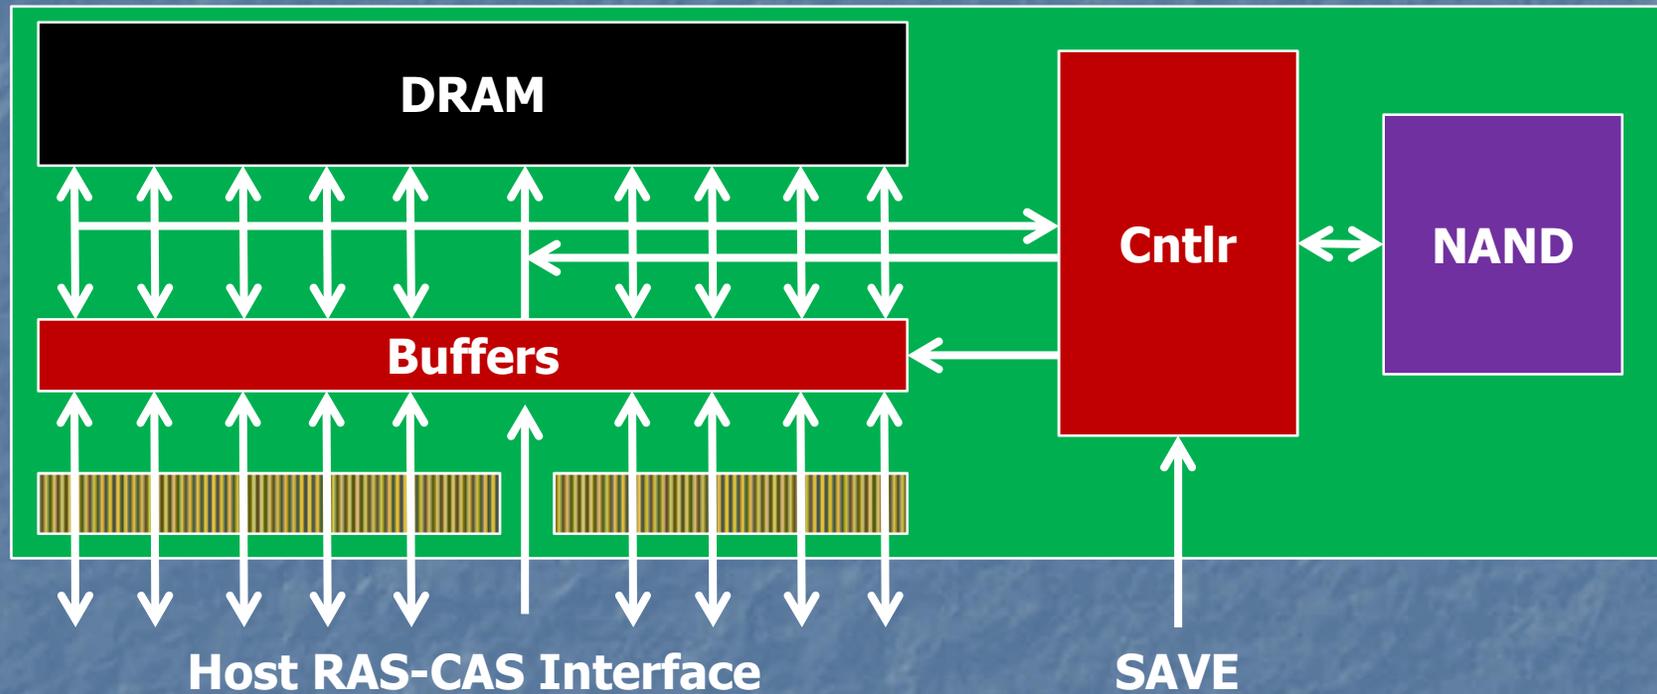
**-N = Persistent DRAM**

**-F = Block accessed NAND Flash**

**-P = Persistent + Block accessed**



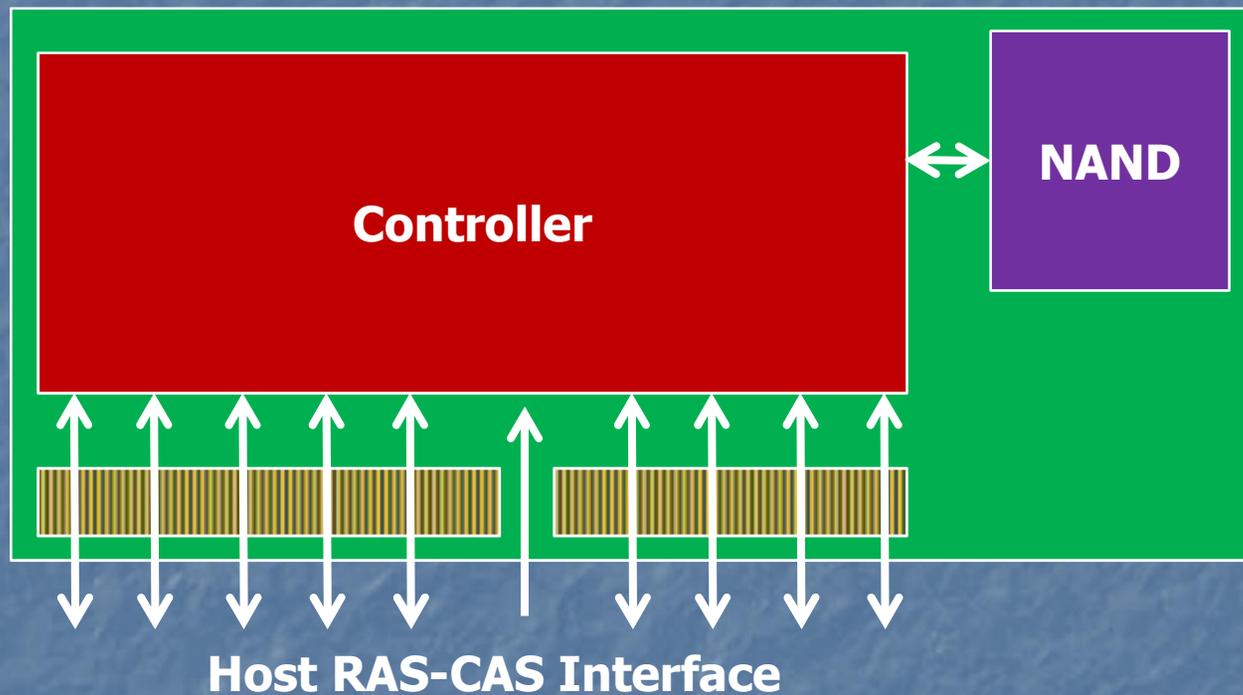
# NVDIMM-N: DRAM Persistence



DRAM accessed at DRAM speeds  
Contents saved to NAND on power fail  
Restored to DRAM when power resumes



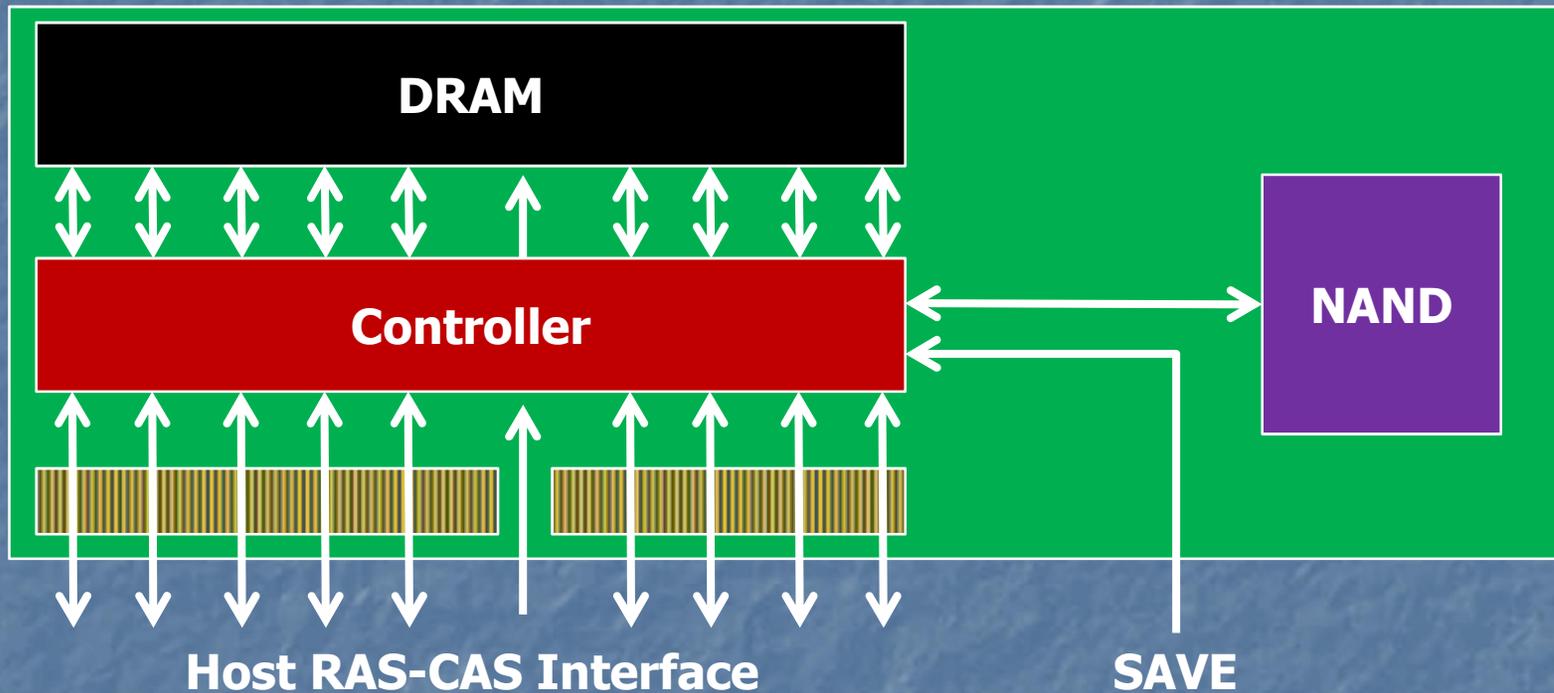
# NVDIMM-F: Block Accessed NAND Flash



No DRAM  
Flash accessed in native block format



# NVDIMM-P: DRAM Persistence & Block



Combined features of -N and -F



# NVDIMM Interface Types

		Function(s)		
		Persistent	Block	Combined
Interface type	Unbuffered	NVUDIMM-N	NVUDIMM-F	NVUDIMM-P
	Registered	NVRDIMM-N	NVRDIMM-F	NVRDIMM-P
	Load Reduced	NVLRDIMM-N	NVLRDIMM-F	NVLRDIMM-P



# Hybrid Enablement

## New signal paths

SAVE pin indicates power fail

12 volts supplied at edge connector

Hybrid controllers appear on I<sup>2</sup>C Bus

## Enhanced SPD Definition

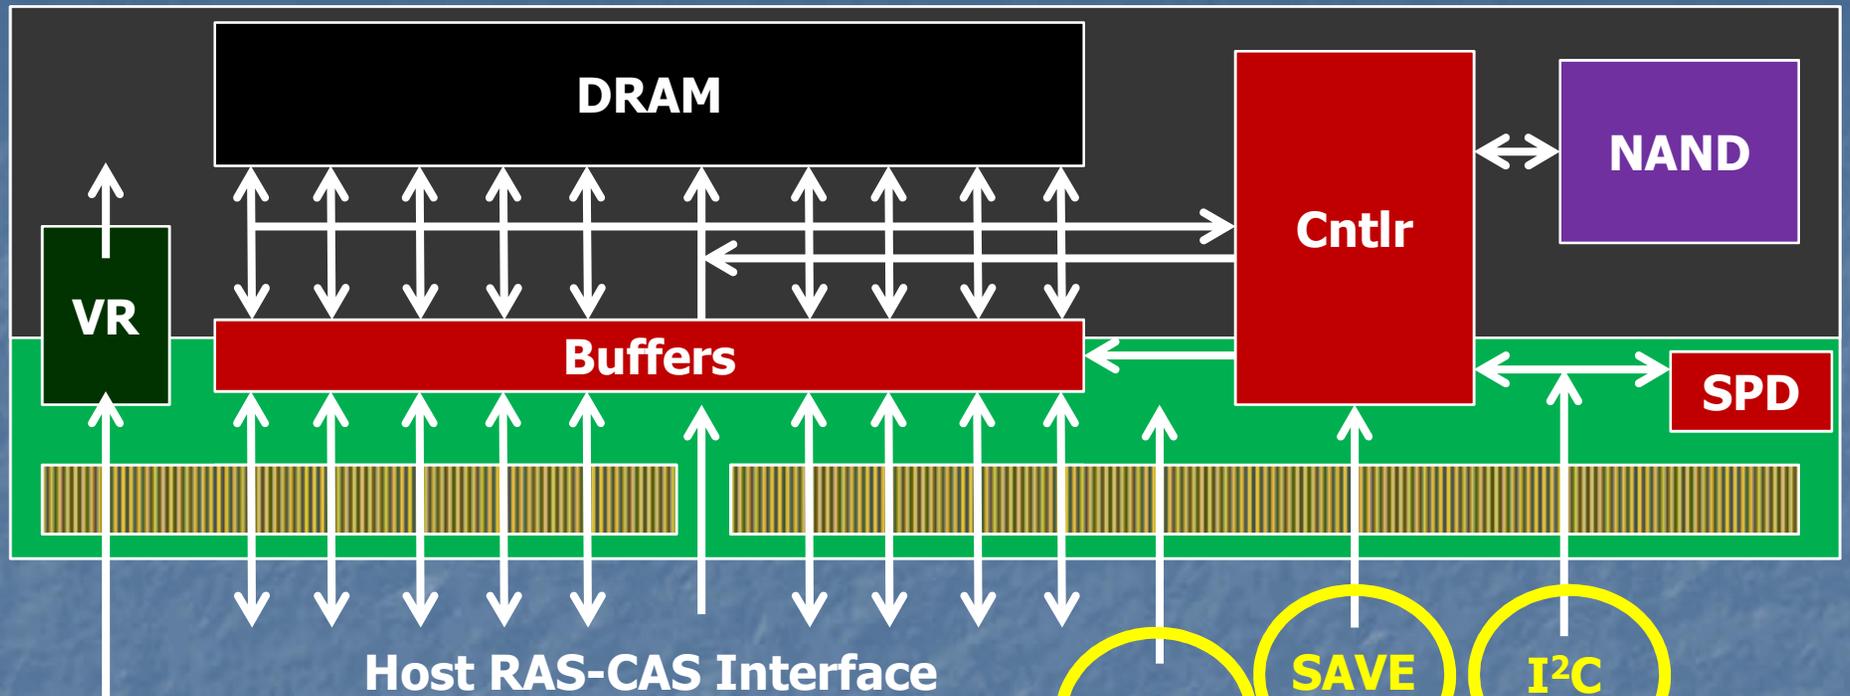
SPD adds new data block for Hybrid

Overlay method extended for Hybrid

## Software Driver Interface

## Standard module label specification





12V

New voltage powers logic during power fail save operation

VDD

VDD powers module during normal op

SAVE

New pin warns that power has failed

I<sup>2</sup>C

I<sup>2</sup>C bus connects to SPD, register, and control logic



## General DRAM definitions

tCK, tRAS, CL,  
ranks, banks, etc

"DDR4-2400P"

**SPD overlay system  
extended to include  
hybrid module data**

## Interface type definitions

LRDIMM,  
register programming,  
etc.

"LRDIMM Interface"

## Hybrid function definitions

NV media types,  
functions, etc.

"NVDIMM"

Extended hybrid data 0

"Persistence"

Extended hybrid data 1

"Block access"

Extended hybrid data 2

"etc"

Extended hybrid data 3



## SPD responsibilities have not changed

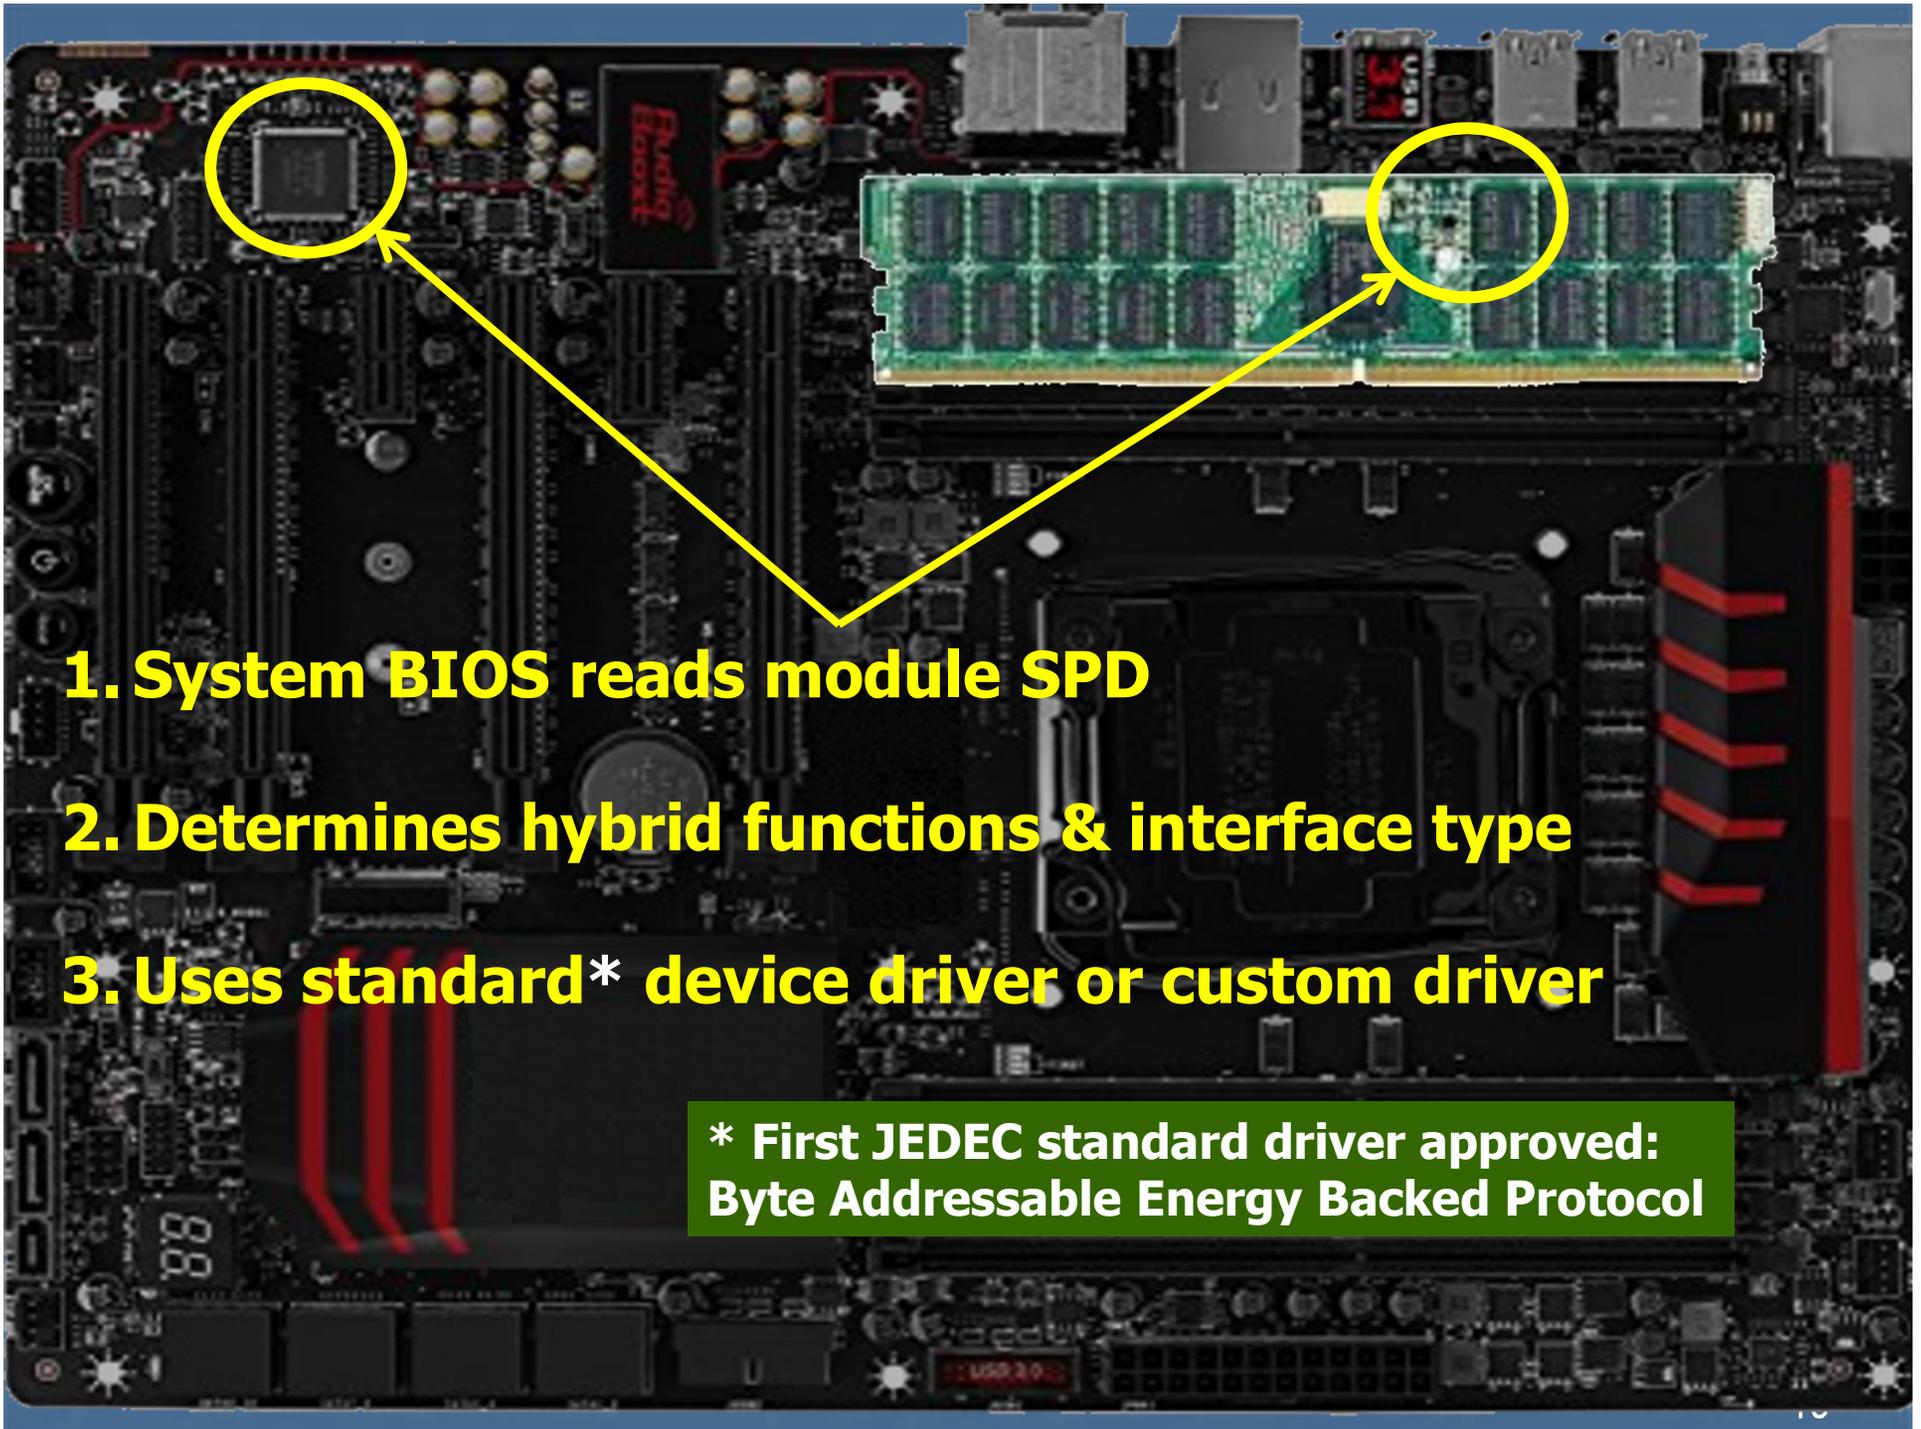
1. Legacy systems must allow newer modules
2. Newer systems must allow legacy modules

Hypothetical SPD Revision Progression Showing Revision Relationships				
Event	UDIMM	RDIMM	LRDIMM	NVDIMM
Initial SPD release	1.0	1.0	1.0	1.0
Addition in RDIMM Annex	1.0	1.1	1.0	1.1
Addition in LRDIMM Annex	1.0	1.1	1.1	1.2
Addition in LRDIMM Annex	1.0	1.1	1.2	1.3
Addition in General Section	1.1	1.2	1.3	1.4
Addition in UDIMM Annex	1.2	1.2	1.3	1.5
Addition in NVDIMM Annex	1.2	1.2	1.3	1.6
Addition in RDIMM and UDIMM Annexes	1.3	1.3	1.3	1.7
Encoding change in LRDIMM Annex	1.3	1.3	2.3	2.7
Addition in LRDIMM Annex	1.3	1.3	2.4	2.8
Encoding change in General Section	2.3	2.3	3.4	3.8
Addition in RDIMM Annex	2.3	2.4	3.4	3.9

## But... SPD revisioning got more complicated

- Module interface type revision dependent on DRAM spec revision
  - Hybrid module revision dependent on module interface type





**1. System BIOS reads module SPD**

**2. Determines hybrid functions & interface type**

**3. Uses standard\* device driver or custom driver**

**\* First JEDEC standard driver approved:  
Byte Addressable Energy Backed Protocol**

# New Standard NVDIMM Module Labels

**gggGB+xxxGB pheRxff Nn4s-wwwaa-mccd-bb**

Persistence  
capacity

Block accessed  
capacity

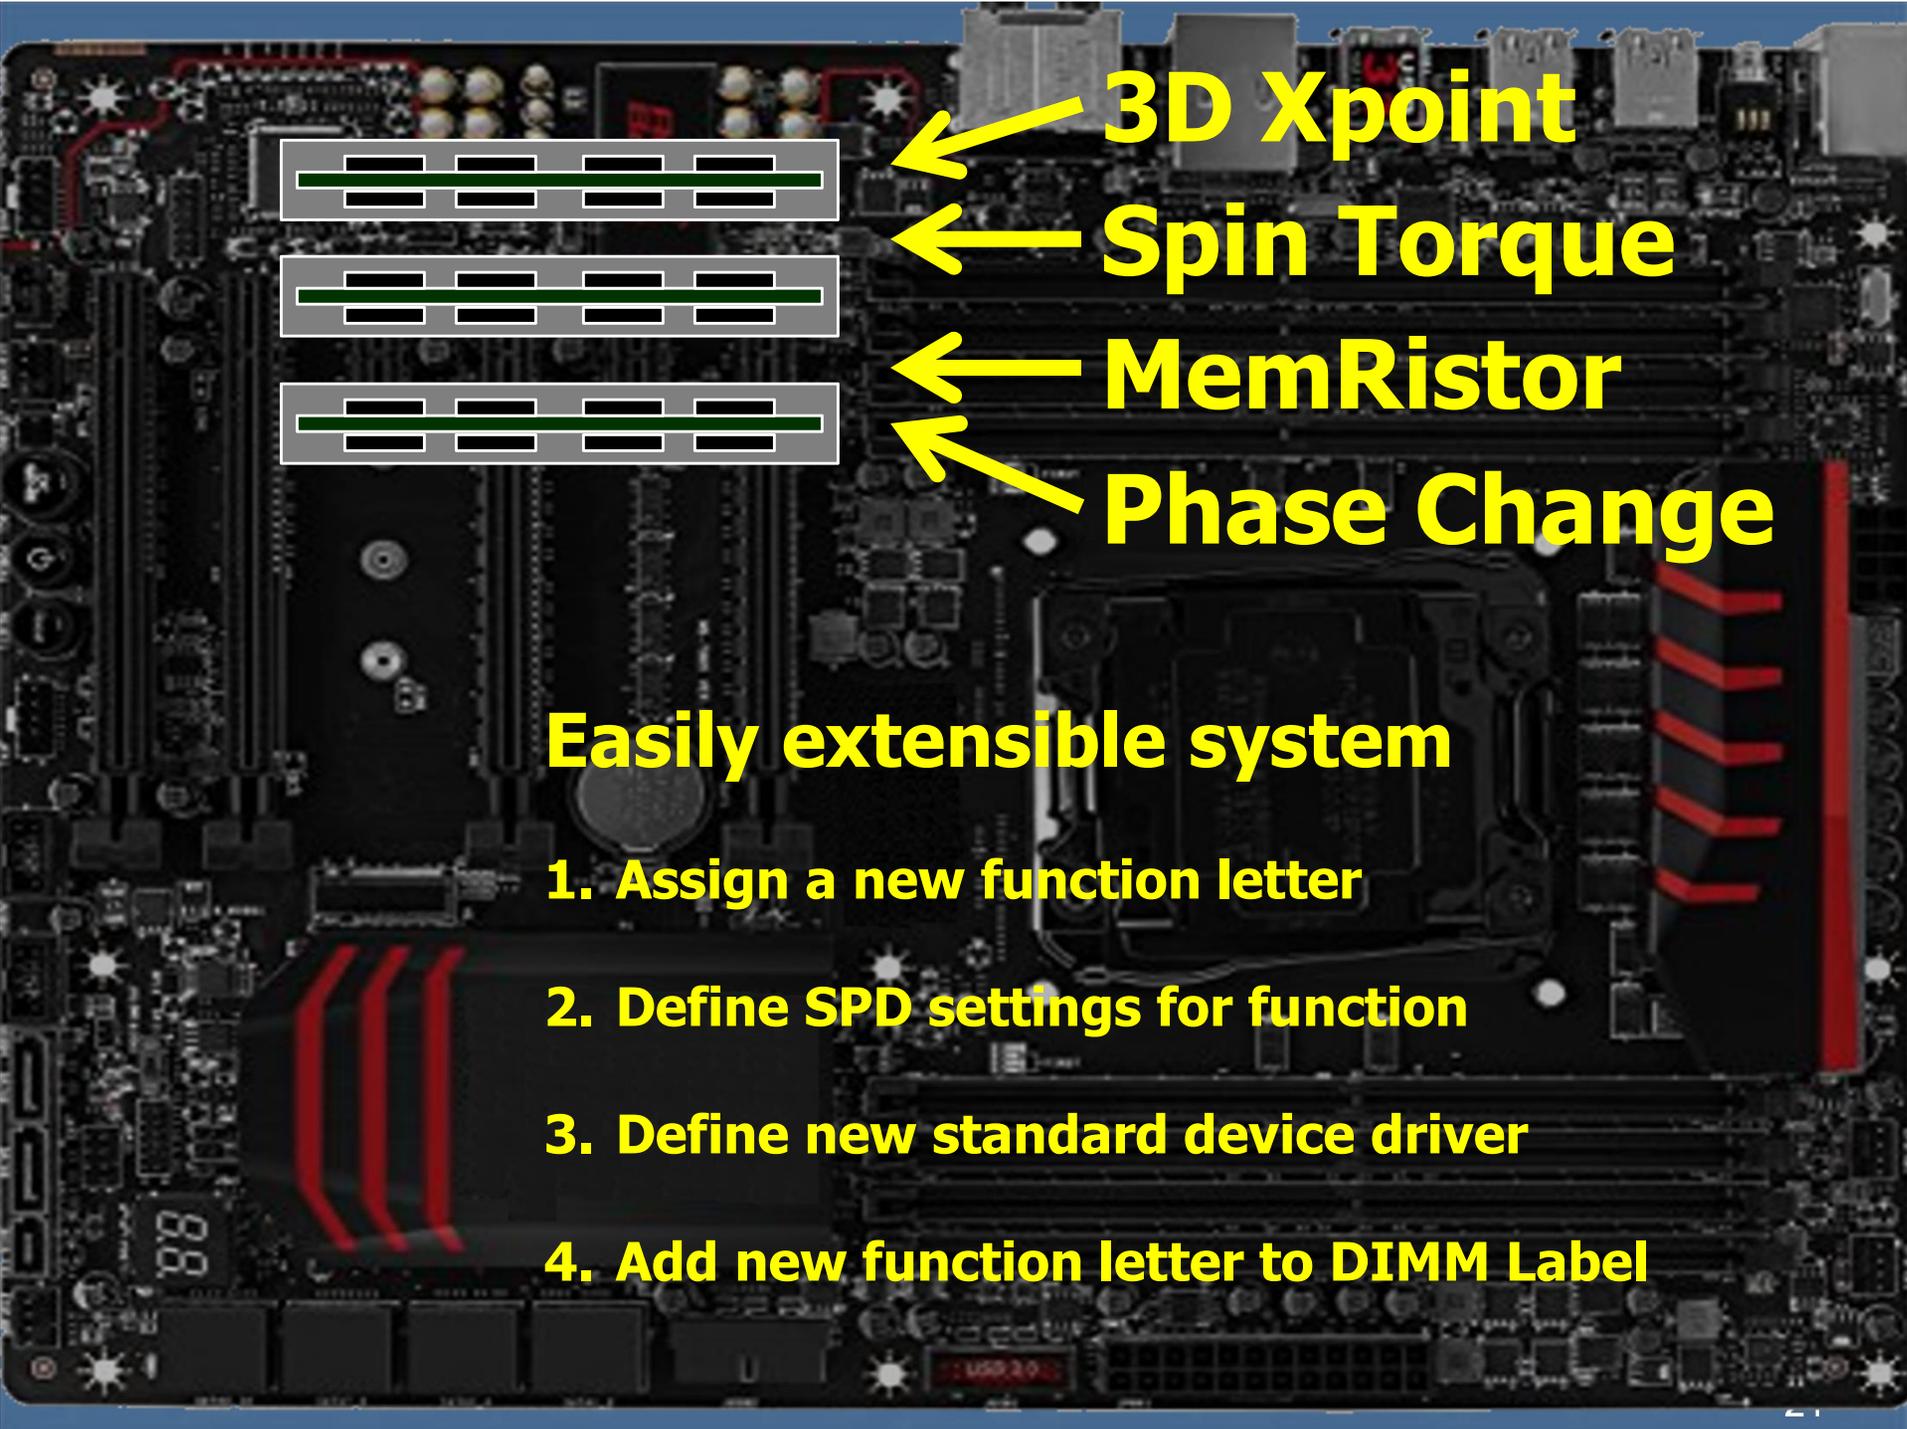
Function type

Interface type

**16GB+32GB 2Rx4 NP4-2133N-RA2-12**

- 16 GB DDR4 SDRAM with persistence
- 32GB NAND Flash mounted as a block oriented device
- NVRDIMM-P (DDR4 RDIMM-compatible interface)
- 2 package ranks per DIMM using SDP DDR4 SDRAMs
- x4 data organization per SDRAM
- DDR4-2133 performance
- Speed grade N: CAS Latency = 14
- Reference design file A revision 2 used for the assembly
- DDR4 SPD revision 1.2





**3D Xpoint**

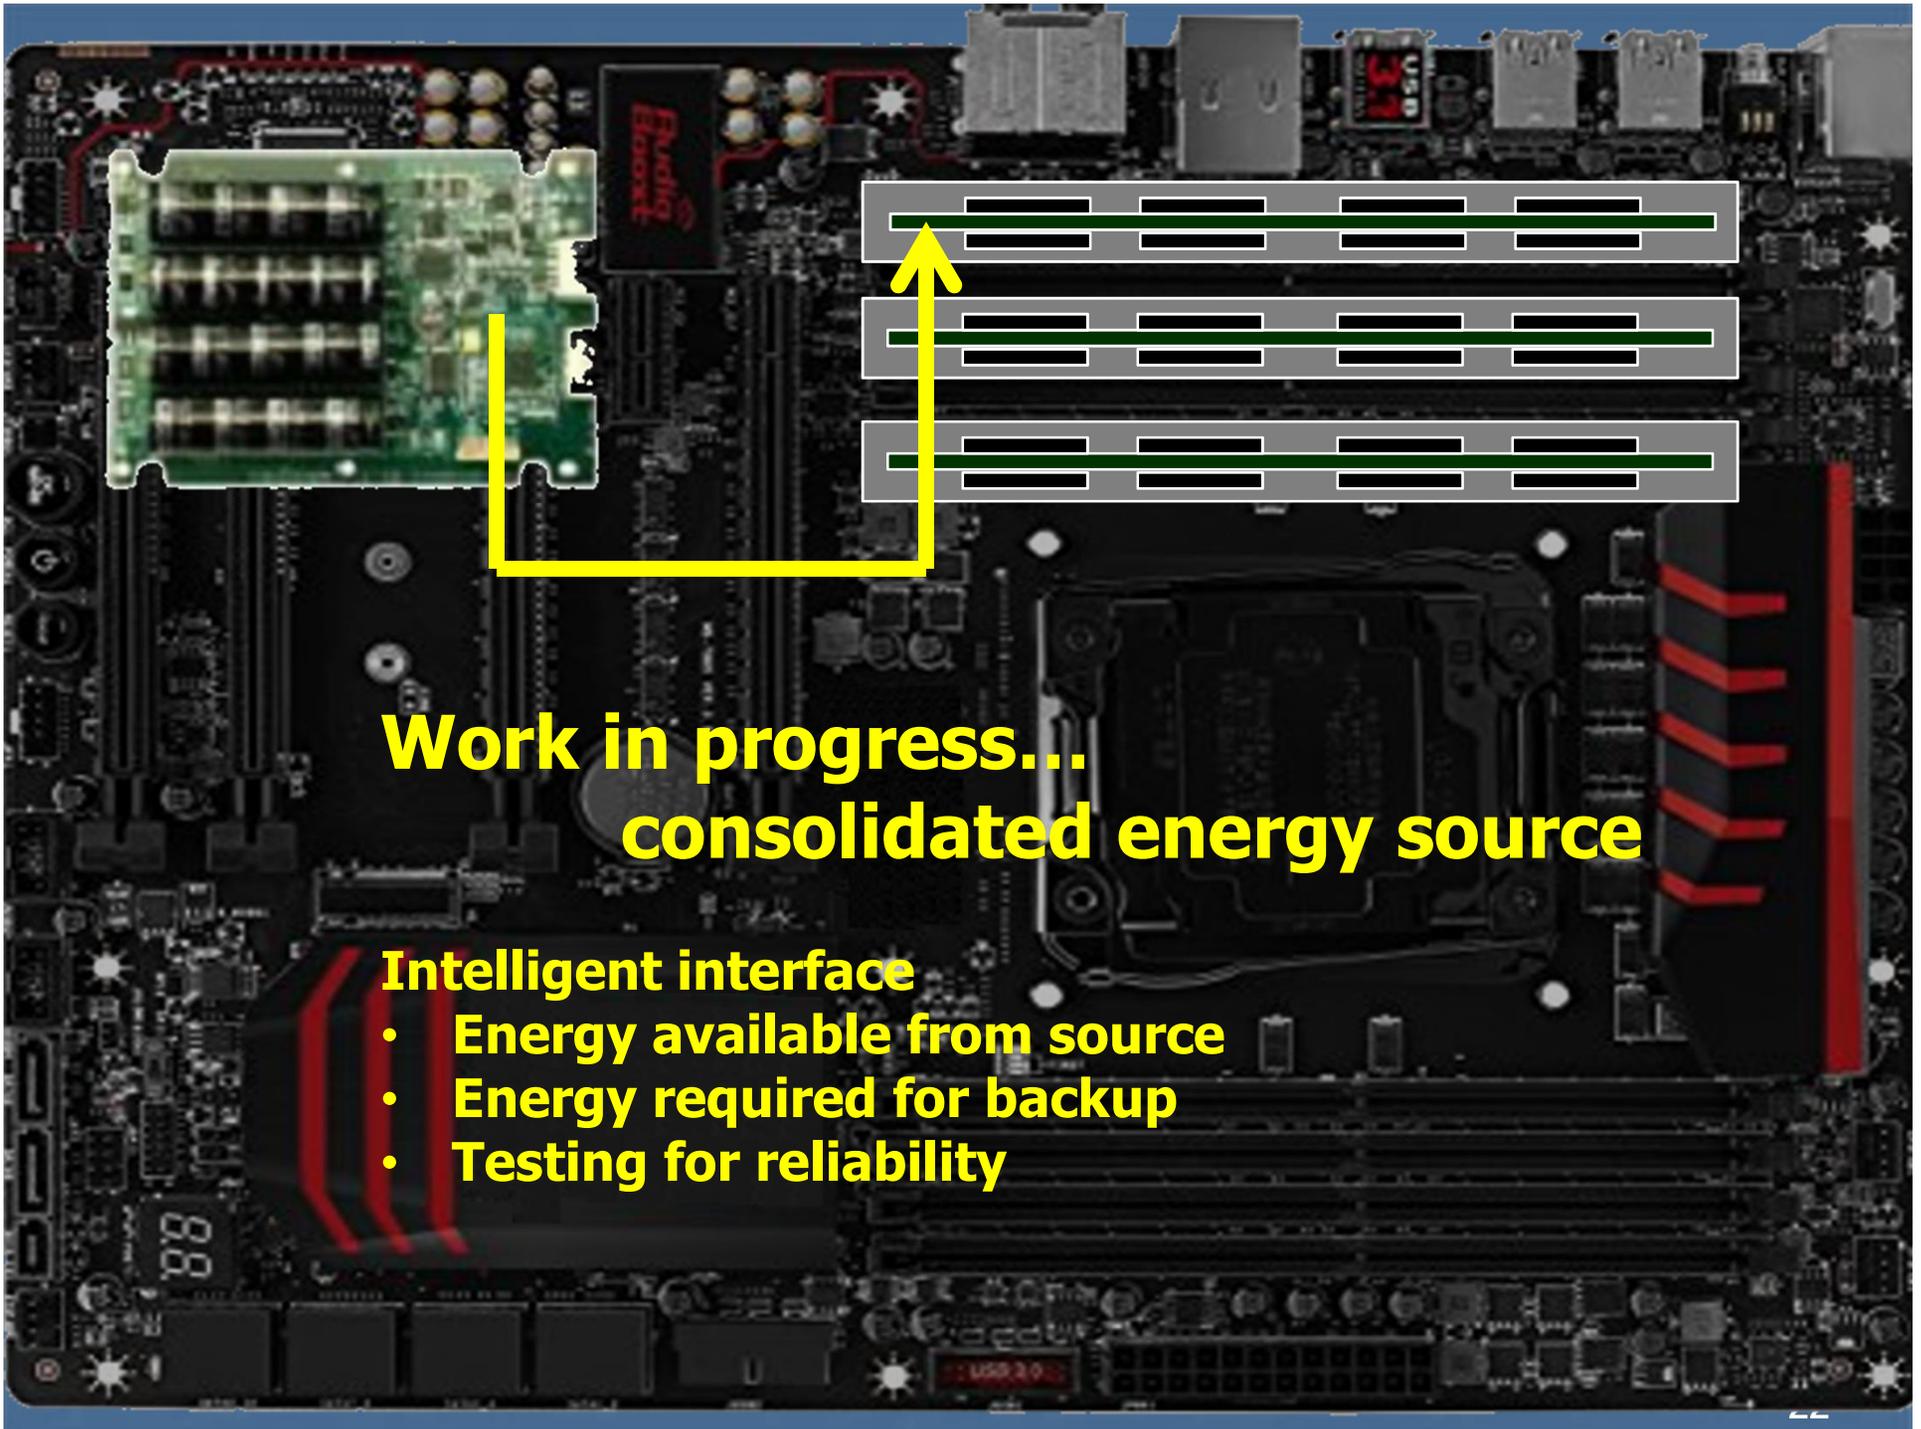
**Spin Torque**

**MemRistor**

**Phase Change**

**Easily extensible system**

- 1. Assign a new function letter**
- 2. Define SPD settings for function**
- 3. Define new standard device driver**
- 4. Add new function letter to DIMM Label**

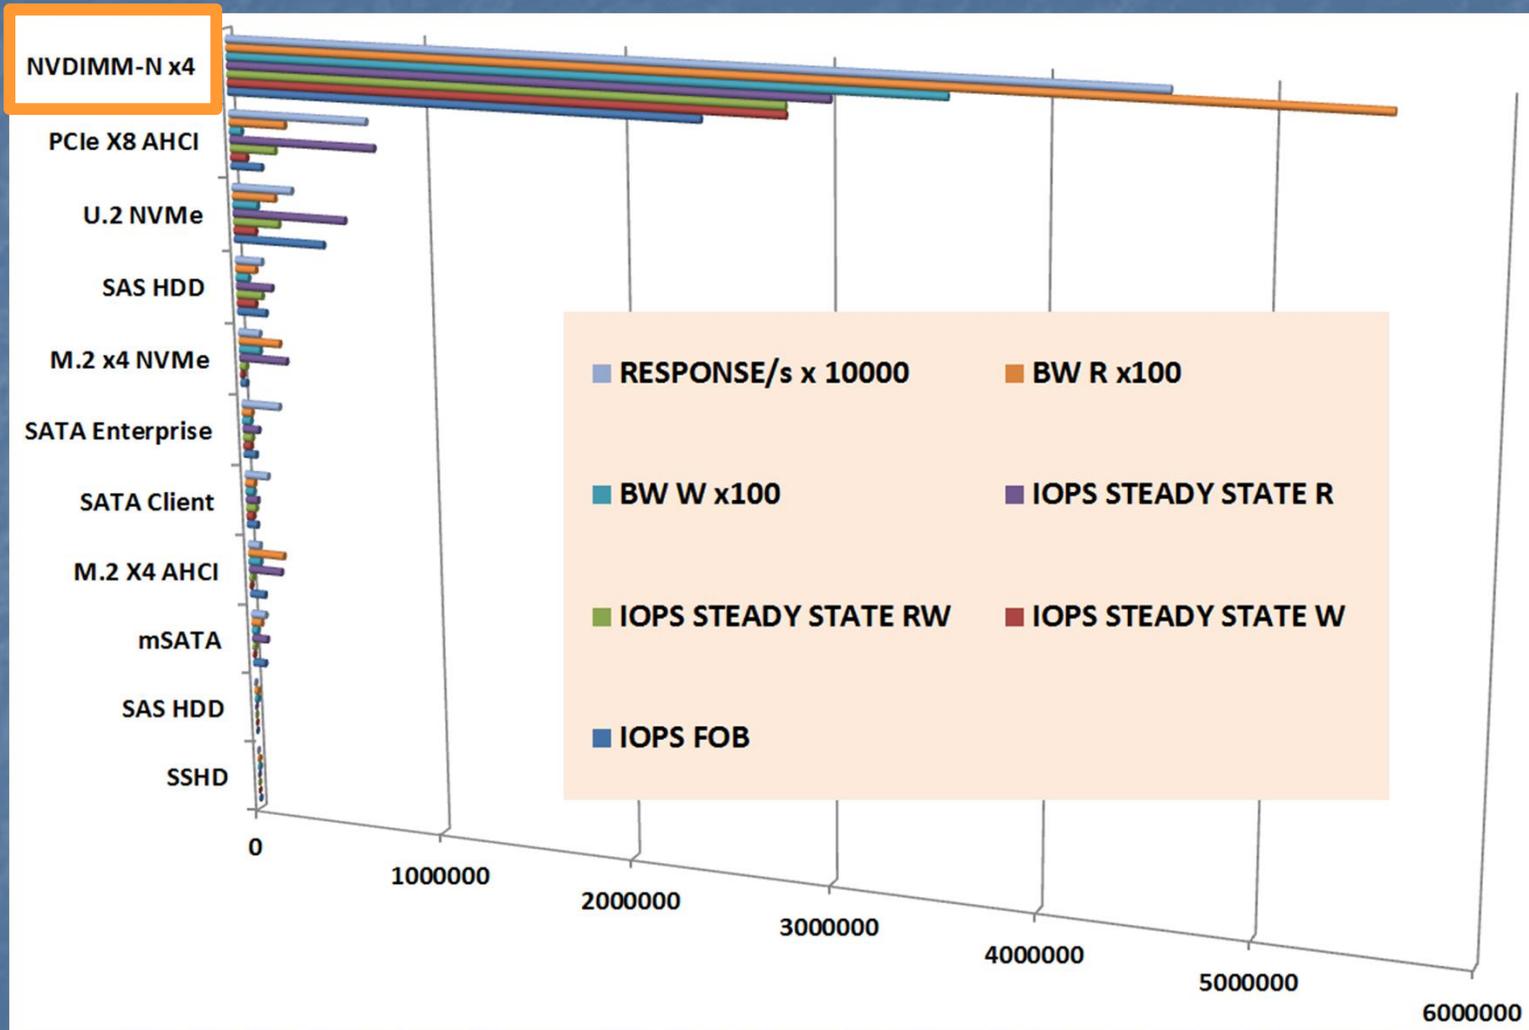


**Work in progress...  
consolidated energy source**

**Intelligent interface**

- Energy available from source
- Energy required for backup
- Testing for reliability

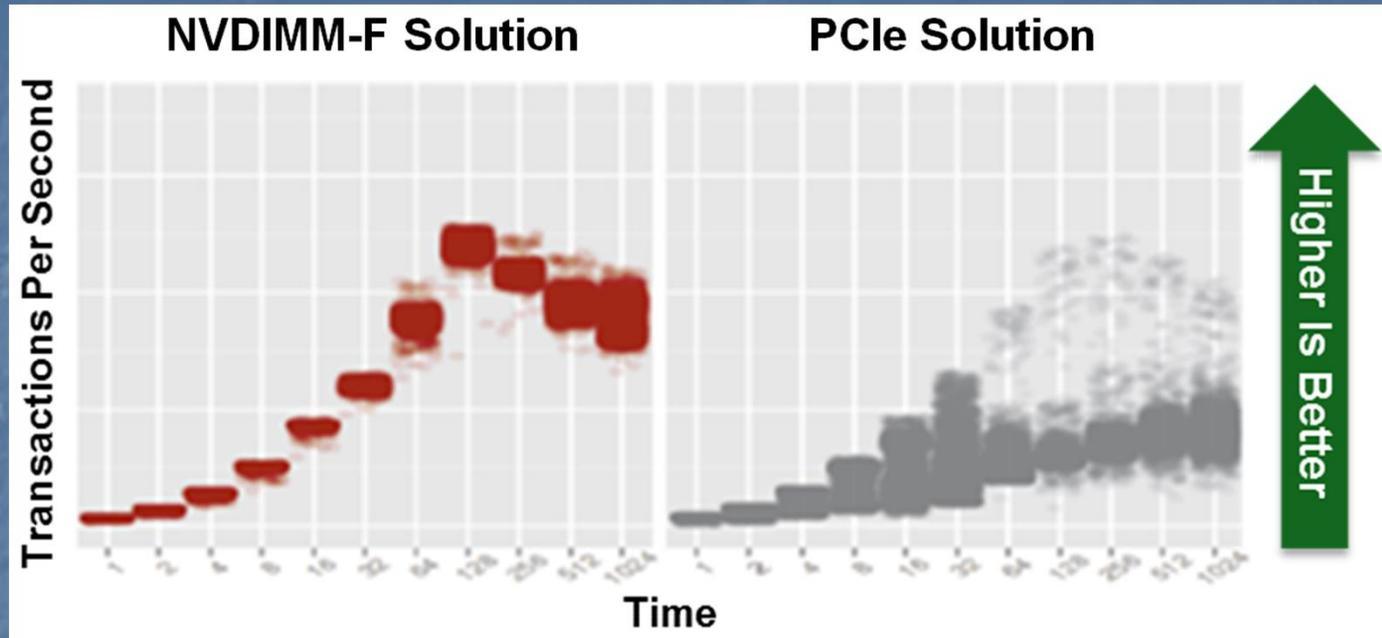
# NVDIMM-N Performance Example



Courtesy: SNIA, Smart, Viking



# NVDIMM-F Performance Example



➤ **3X** Increase In MySQL Transactions Per Second



➤ **4X** Faster Response Time

Courtesy: Diablo Technologies



# Summary

Hybrid modules protocol based on DRAM

Acts like UDIMM, RDIMM, or LRDIMM

Hybrid & DRAM modules can share the bus

NVDIMM types defined so far

-N = persistence

-F = block access

-P = -N + -F

New power and signals enable HW

New SPD structure enables SW

Standard device drivers enable usage

New DIMM labels assist end user selection



**JEDEC**

**Thank you**

Bill Gervasi

[bilge@discobolusdesigns.com](mailto:bilge@discobolusdesigns.com)



**DISCOBOLUS**  
**DESIGNS**