

MEMCON08
DENALI SOFTWARE, INC.

JULY 21-24 • SAN JOSE, CA



NVMHCI - Coming Out From Behind the Curtain

By Nathan Obr
storage driver developer
Microsoft

Microsoft

Microsoft

Non-Volatile Memory Host Controller Interface (NVMHCI)

- In the next 25 minutes ...
 - Overview
 - Committee and Standard
 - Architecture
 - High Level Features

NVMHCI Overview

- New Development
 - NVM has become cost effective in PC relevant sizes
- Existing Problem
 - Storage is a performance bottleneck in the PC
- Opportunity
 - Mainstream NVM into PC architecture in a performant and standardized form
- Solution
 - Standardize an NVM HCI interface to promote adoption

NVMHCI Committee and Standard

- Formed May 2007
- NVMHCI 1.0 was published on April 14, 2008.
- Leveraged AHCI semantics, concepts and experience
- Currently consists of 35+ members



NVIDIA




FRESCO LOGIC



AMD

Smarter Choice



OSR

SanDisk



MARVELL

DELL

Skymedi

GENESYS LOGIC

HITACHI

Inspire the Next



IP CORES



SiliconMotion

PHISON

Knows What You Need



JMicron



Silicon Integrated Systems Corp.


Seagate



Microsoft


Microsoft

Storage Bus Architecture




Legend
CI = Command Interface
DC = Disk Controller
HCI = Host Controller Interface

Storage Bus Architecture with SSD




Legend
NC = NVM Controller
NVMC = NVM Channel

NVMHCI Architecture



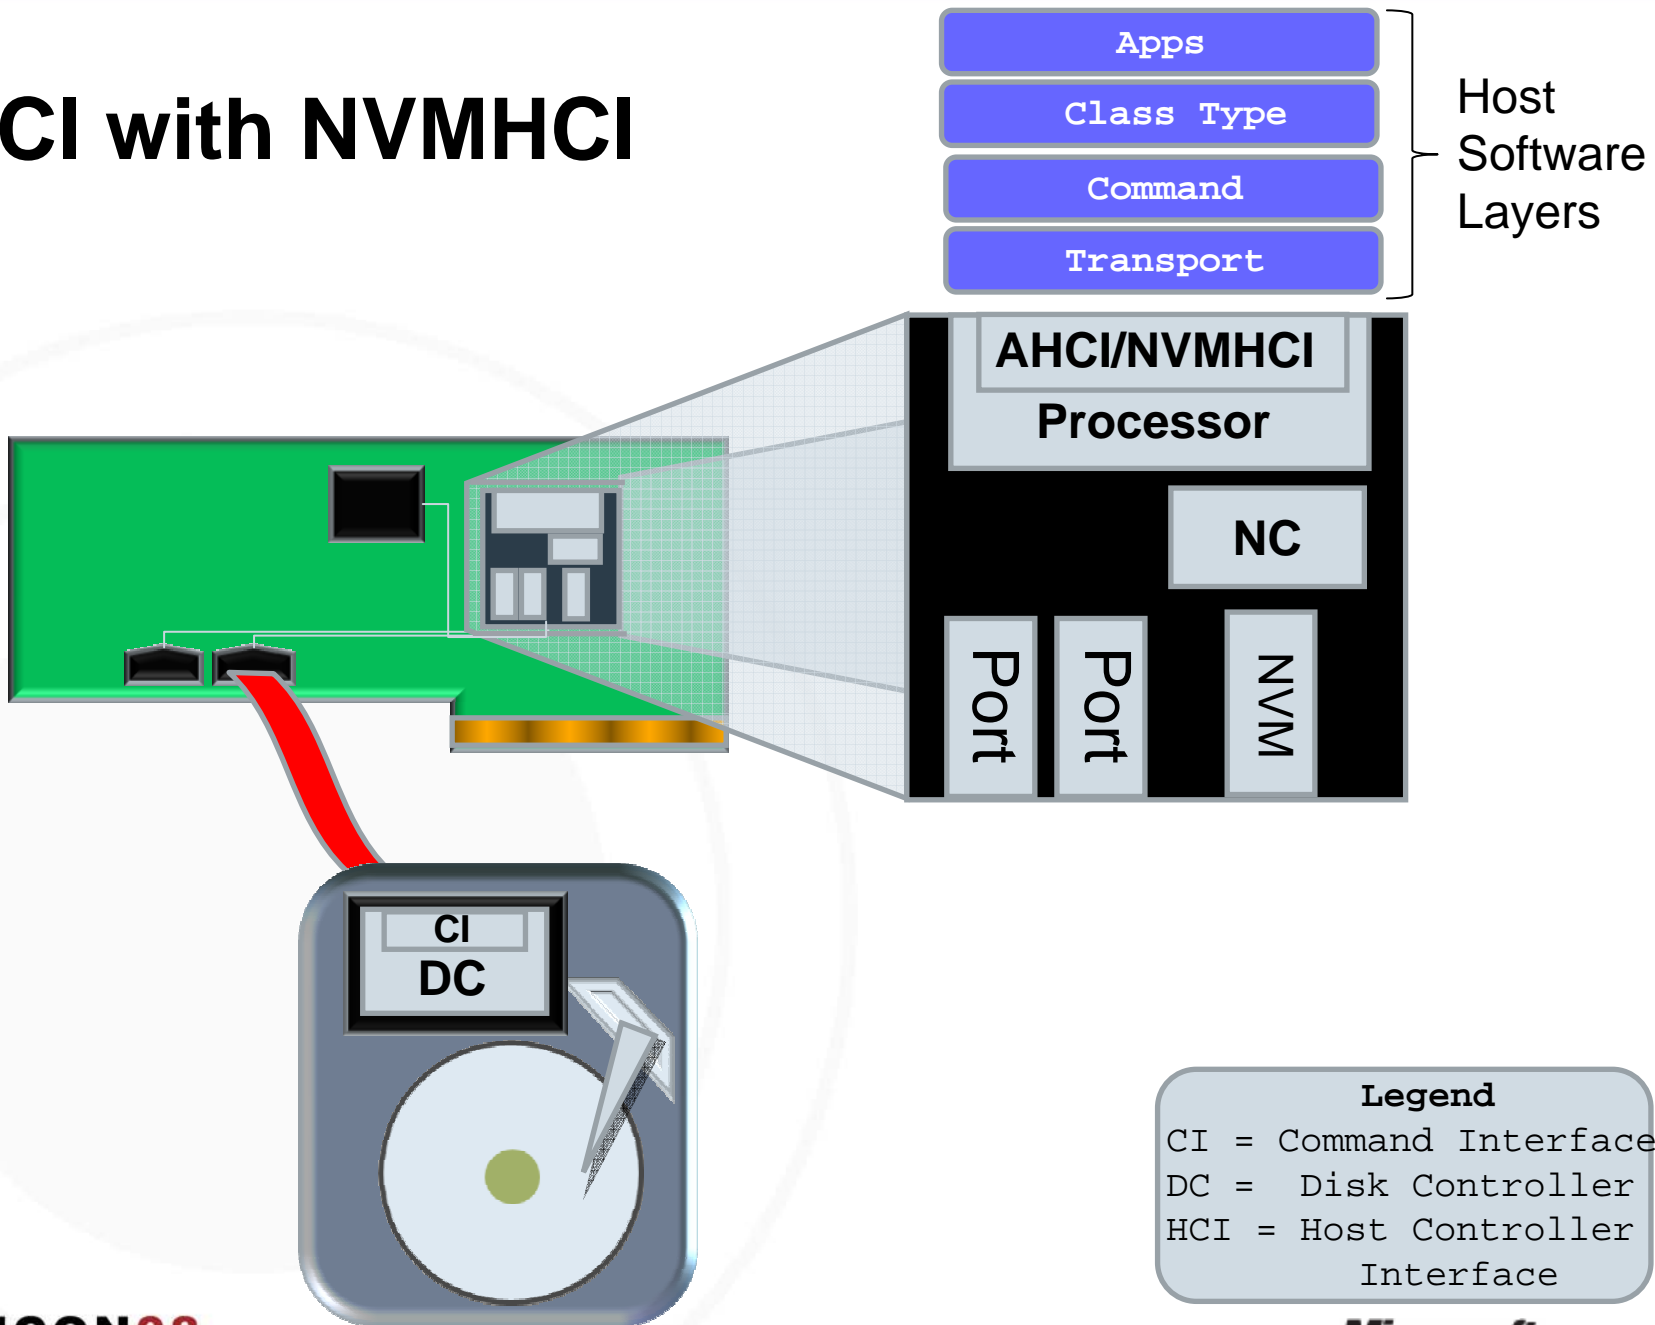
Legend
NC = NVM Controller
NVMC = NVM Channel

Storage Bus vs. NVMeHCI Architecture



Legend
↔ = Equivalent Purpose

AHCI with NVMHCI



Legend
CI = Command Interface
DC = Disk Controller
HCI = Host Controller Interface



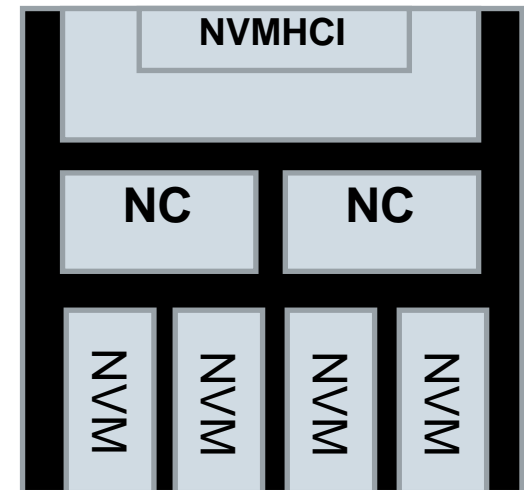
NVMHCI Key Benefits

NVMHCI standardizes an NVM host controller interface that:

- prevents legacy device and bus limitations
- minimizes controller complexity
- enables features only available on NVM controllers
- enables caching solutions

... prevents legacy device and bus limitations

- Asynchronous notification to the host
 - Enables a controller to request an action by the host on its behalf
 - Multiple simultaneous commands
 - Preserves parallelism of NVM controller designs
- Out of order data transfers
 - Improves throughput efficiency by using NAND as it is ready as opposed to maintaining the DMA engine's traditional order



... minimizes controller complexity

- Inseparable Host controller and NVM controller
 - The controller can be used to manage the storage in a consistent manner over the course of the media's life
- Focused operation set ...

Enumeration	Identify
Configuration	Get Features, Set Features
Health Monitoring	Get Status
IO	Read, Write, Flush
Management	Data Set Management

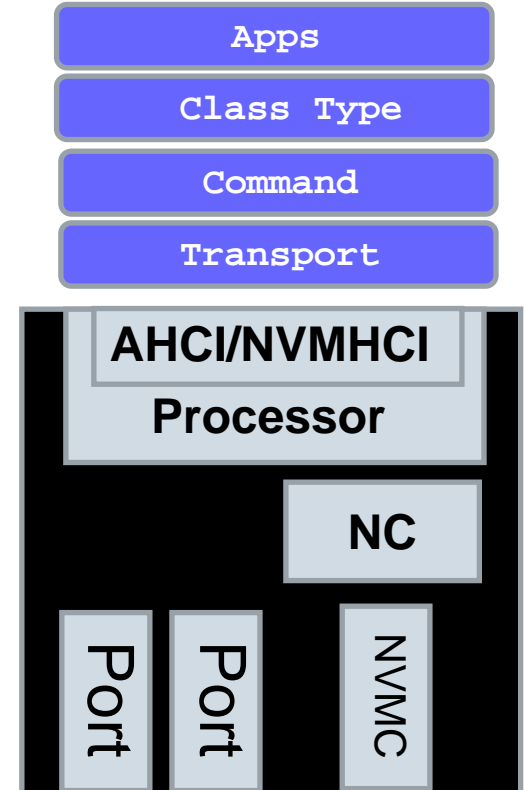
... enables features only available on NVM controllers

- Operations defined in a large, extensible, segmented structure
 - Timeout
 - Workload Management
 - Prioritization
- Metadata transfers
 - Provides metadata read/write access to the host. Metadata access can even be done during IO operations.

Command
Address Low
Address High
Transfer Count
Parameters
Attributes
Reserved
Reserved

...enables caching solutions

- BIOS support
 - Leverages an NVMHCI boot cache
 - Prevents boot failures using safe mode
 - Cache usage control for OEMs
- Device Layout Types
 - prevents cache segments from being used as permanent storage and vice versa.
- NVMHCI ports under an AHCI controller
 - used as cache for a SATA HDD.



But wait, there's more!

- Multiple ports
- Enumeration efficiency
- Index Data/Pair
- Data Set Management
- Logical, Physical, Page alignment
- Configuration Persistence
- Metadata size
- Linear processing
- Power Management
- Cache Map Table Formats
- Allocation and Optimum Transfer size

Next Steps

- Get the spec
 - <http://www.intel.com/standards/nvmhci/index.htm>
- See what NVMHCI 1.0 can do for you
- Join the committee