



# NVMe OFA Open Source Windows Driver

Kwok Kong

Sr Director of Software Engineering

PMC-Sierra



# Agenda

- Status Update
- Driver Architecture
- Driver Features
- Future

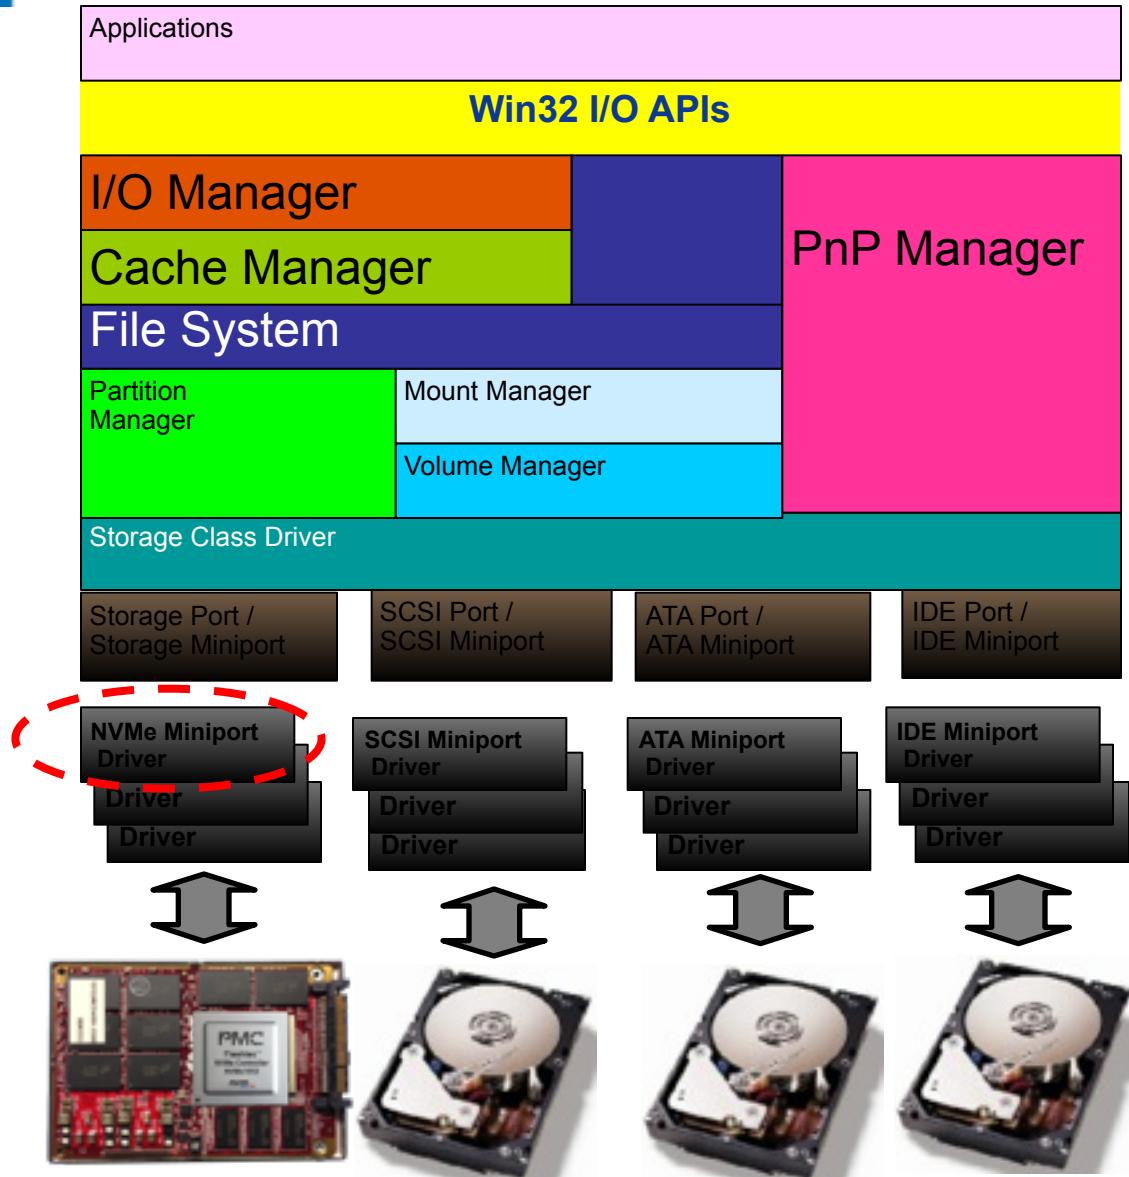


# Status Update

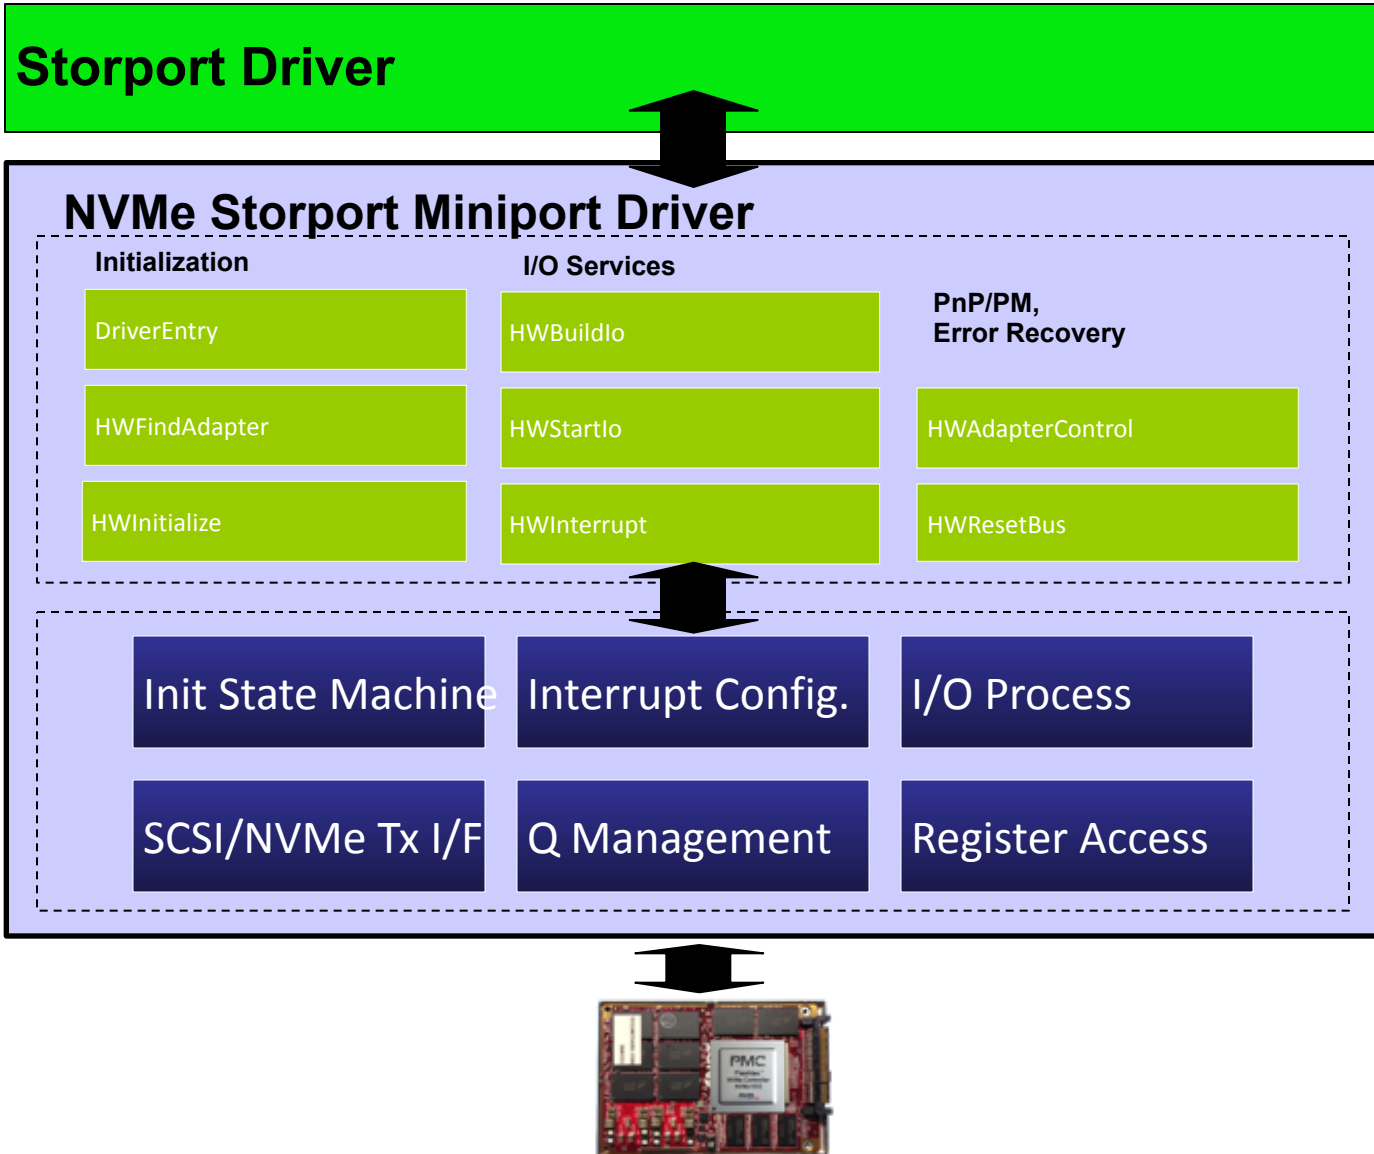
- Release 1**
  - Q2 2012 (released)
  - 64-bit support on Windows\* 7, Windows\* Server 2008 R2
  - Mandatory features
- Release 1.1**
  - Q4 2012 (released)
  - Added 64-bit support Windows\* 8
  - Public IOCTLs and Windows\* 8 Storport updates
- Release 1.2**
  - Q2 2013 (released)
  - Added 64-bit support on Windows\* Server 2012
  - Signed executable drivers
- Release 1.3**
  - Q2 2014 (released)
  - Added 32-bit support on all supported OS versions
  - Hibernation and better MUMA group support

Four major releases of the Windows\* OFA community driver since 2012  
Code contributions from Huawei, PMC, Intel, LSI, and SanDisk

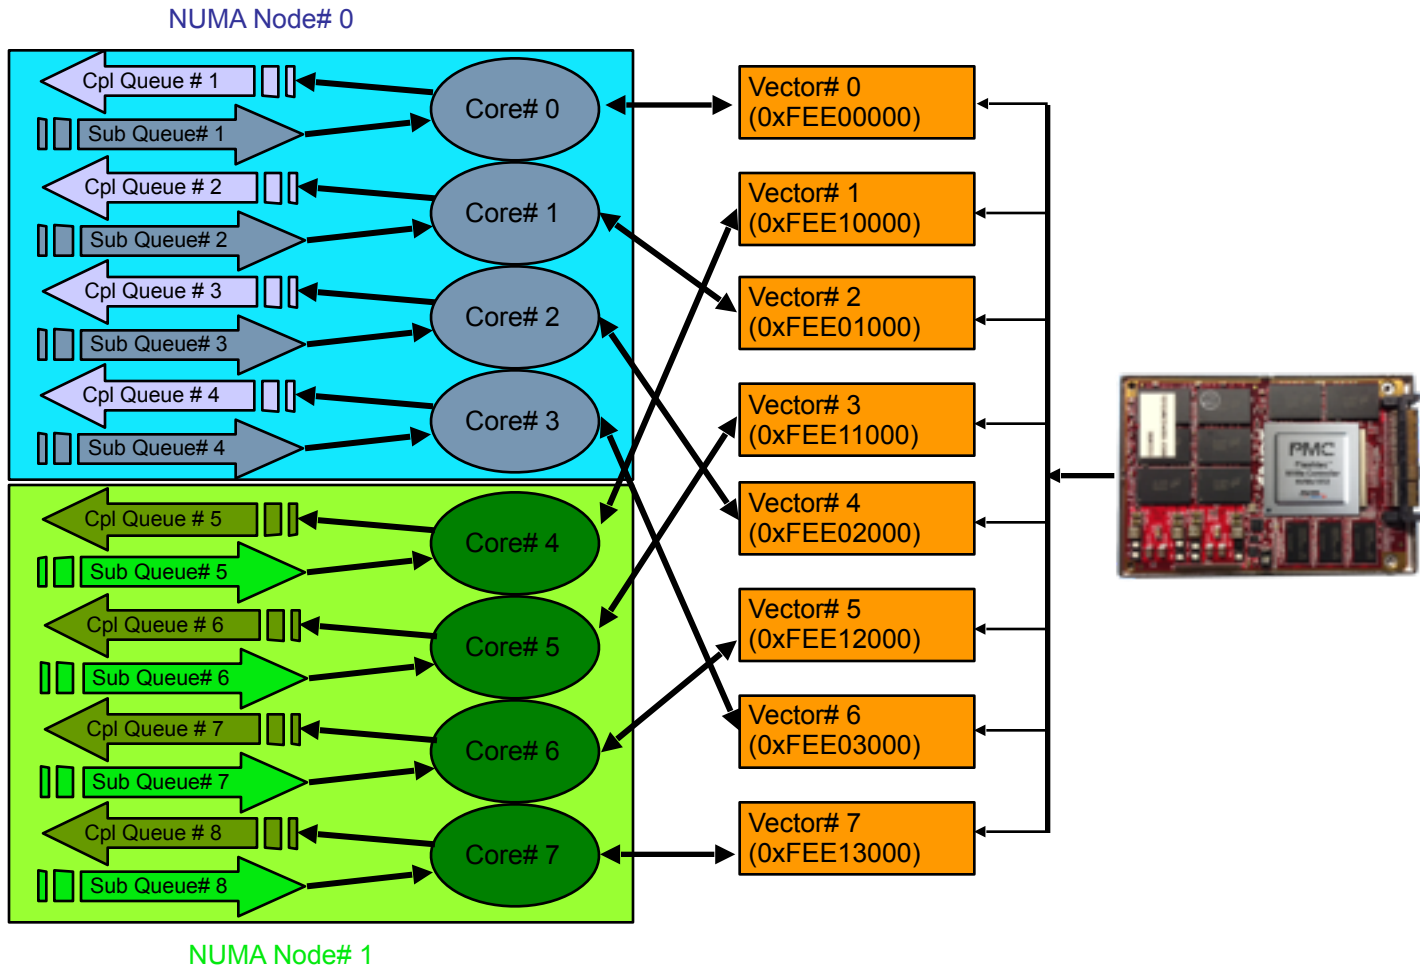
# Windows Storage Architecture



# Driver Architecture



# Queue-Core-Vector Mappings





# System Features Overview

Features	Supported
NVMe Specification	1.0e
Windows Versions (32-bit and 64-bit)	7, 8, Server 2008R2, Server 2012
NUMA Optimized Queues and Memory	Dedicated IO queues per CPU core Single Admin Queue for all CPU cores Queue memory allocated local to NUMA Node
Interrupt	MSI-X (Vectors mapped to NUMA optimized IO Queue Pairs) MSI INTX
Queue Arbitration / Priority	Round Robin only 1:1 mapping between Submission and Completion queues
Pass Through	Yes (with DeviceIoControl())
Multiple Namespaces	Up to 16, Mapped to Bus 0, Tgt 0, LUN 0-15
Registry Parameters	Name Space, Max Transfer Size, Admin Queue Size, IO Queue size, Interrupt Coalescing Time/Entries)
End-to-End Protection	No



# Future Features

- 1.4 Release (end of 2014)
  - WHQL Certification test
  - Driver Trace feature
  - Hot Plug / Swap support
  - NVM Format Enhancement
  
- 2015 and Beyond
  - NVMe 1.1b/1.2 features
  - T10 end-to-end protection

Need your contribution to write the future





# OFA NVMe Driver Working Group

- Founding Companies
  - PMC – Chairperson, Code maintainer
  - Intel
  - Avago (LSI, SandForce)
  
- To Contribute
  - Join the mailing list
  - Email your patch to the WG mailing list
  - Code is checked in if approved by two out of three founding companies



# Useful Information

- Join the Driver Mailing List to Contribute
  - <http://lists.openfabrics.org/cgi-bin/mailman/listinfo/nvmewin>
- Development Tools Description
  - <https://www.openfabrics.org/index.php/developer-tools/windows-tools.html>
- Driver Source Code
  - <http://svn.openfabrics.org/svnrepo/nvmewin>