



NVMe VMware vmklinux Driver

Kwok Kong

Sr Director of Software Engineering

PMC-Sierra



Flash Memory Summit Agenda

- Status Update
- Driver Architecture
- Driver Features
- Future

Status Update

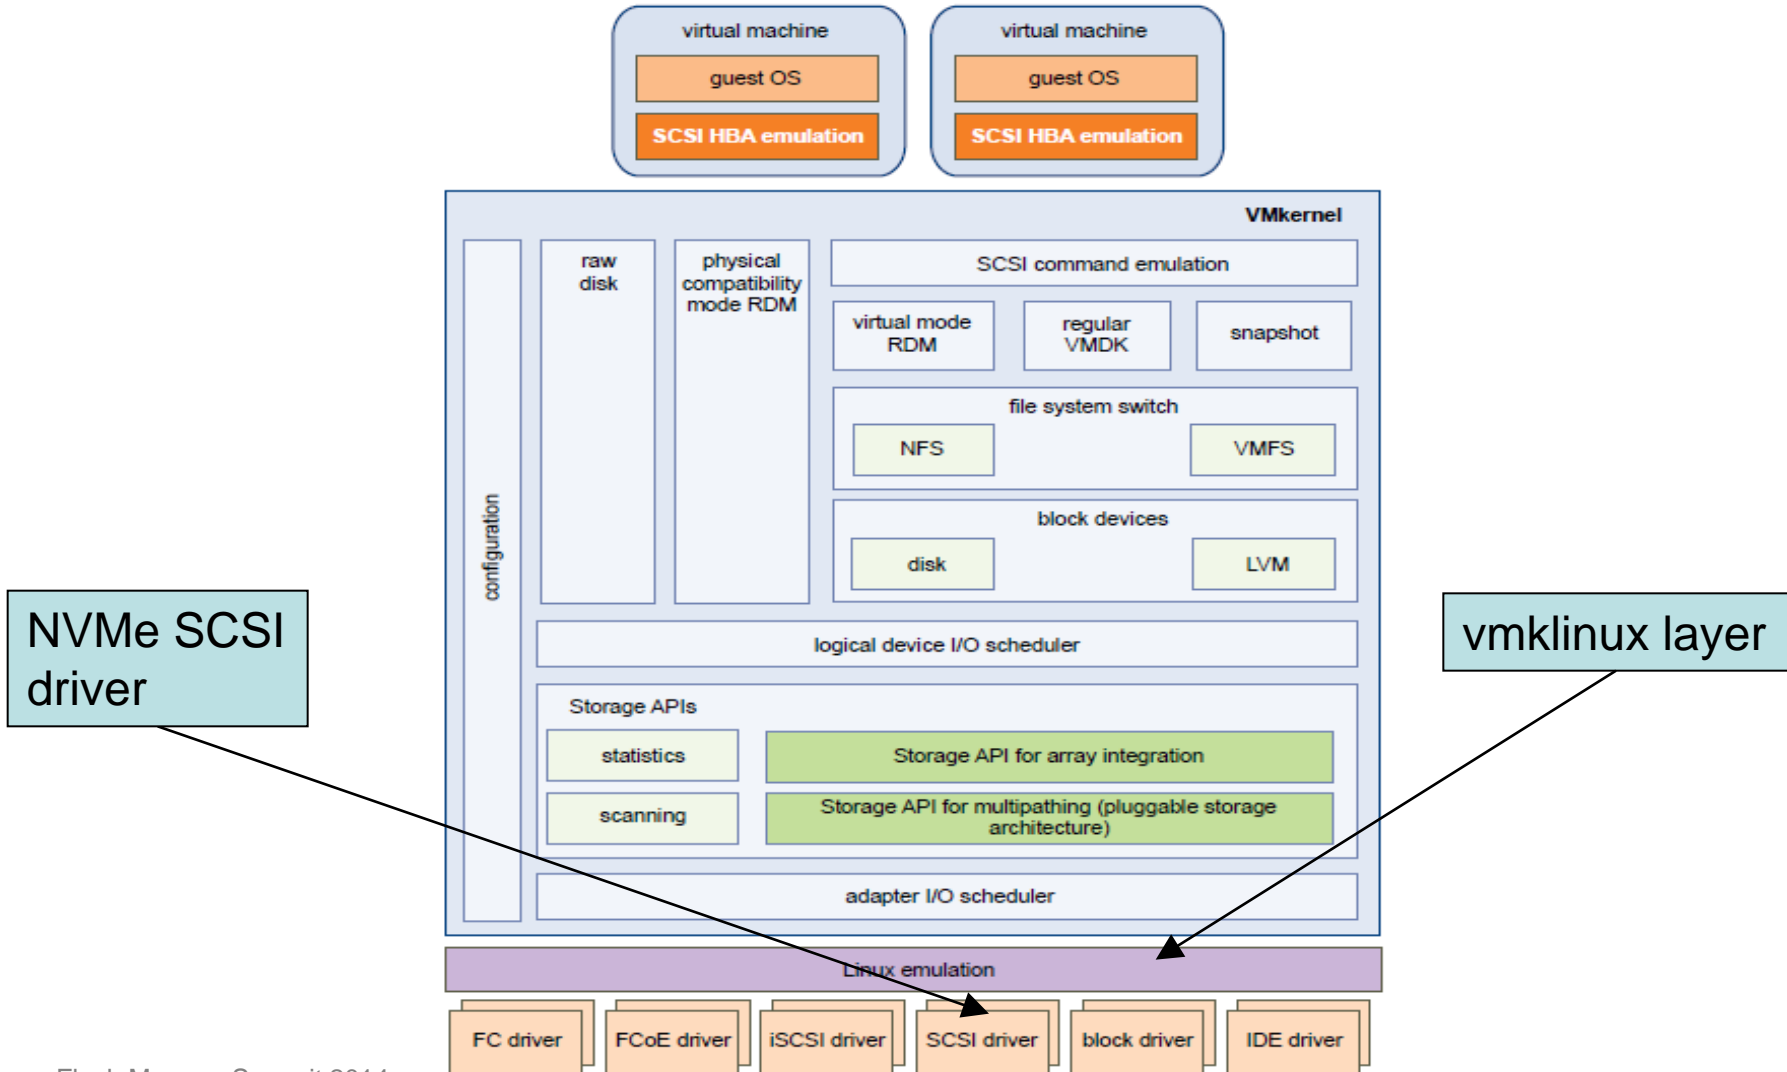
Release 1.0

- Q2 2014 (released)
- Supports ESXi vmkernel 5.0 (beta and release)*
- vmklinux interface *
- SCSI device driver
- Passed 5.0 certification process for local storage with SSD based on PMC NVMe controller

First release

Joint development between VMware and PMC started in Q3 2013

VMware Storage Architecture



Driver Architecture

Vmkernel SCSI Layer



NVMe SCSI Driver

Initialization(pci probe)

ScsiAddHost

Admin queue setup and
Nvme hw initialization

Interrupt setup

IO queues configuration

Scsi scan device

I/O Services

NvmeScsiCommand

NvmeScsiAbort

NvmeQueue_ProcessCq
(processing IO completion)

PnP/PM,Error Recovery

Timeout mechanism

shutdown

Char device management

passthrough

reset

Hot add/remove NS



System Features Overview

Features	Supported
NVM Express version	1.0e
ESXi versions	5.0 (beta and release)
NUMA Optimized Queues	Dedicated IO queues per NUMA node Single Admin Queue for all CPU cores
Interrupt	MSI-X only(coalescing not supported)
Queue Arbitration / Priority	Round Robin only 1:1 mapping between Submission and Completion queues
Pass Through	Yes(both Admin and IO)
Multiple Namespaces	Up to 8, Mapped to Bus 0, Tgt 0, LUN 0-7
Registry Parameters	Name Space, Max Transfer Size, Admin Queue Size, IO Queue size
End-to-End Protection	No
NVMe reset	Yes (using ioctl)



Future Features

Need your contribution to write the future



Useful Information

- Register with Sourceforge
 - <https://sourceforge.net/user/registration>
- Subscribe to the project update
 - <http://sourceforge.net/projects/nvmevmklinuxdriver/>
- Driver Source Code
 - <http://nvmevmklinuxdriver.sourceforge.net/>

Thank You!



See us in booth # 416

www.pmcs.com

blog.pmcs.com [@pmcsierra](https://twitter.com/pmcsierra)