



Flash Memory Summit

Using **Open Channel SSDs** to Improve QoS for Datacenter Applications

Xueshi Yang, PhD
CEO, **Shannon Systems**
August 2018



WHY

HOW

WHAT

Open Channel SSDs for the datacenter

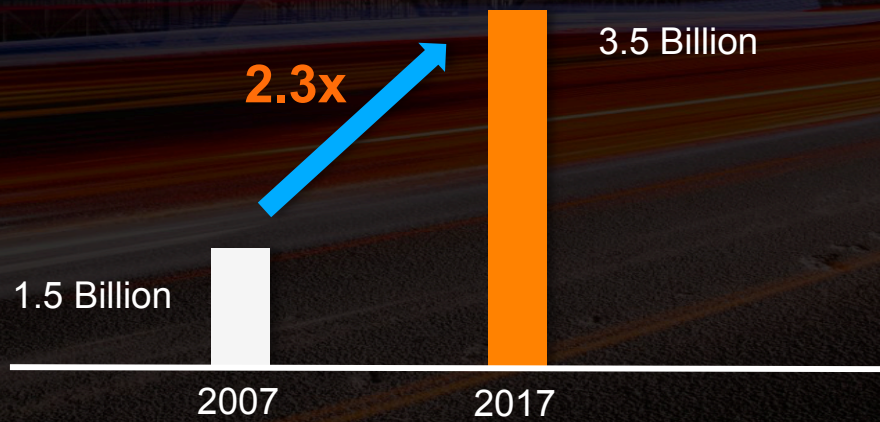
Transactions, the history





Transactions, today

Transactions: Volume



2017 天猫双十一

天猫双十一
GLOBAL SHOPPING FESTIVAL 2017

Transactions: Velocity

2017.11.11 00:03:01

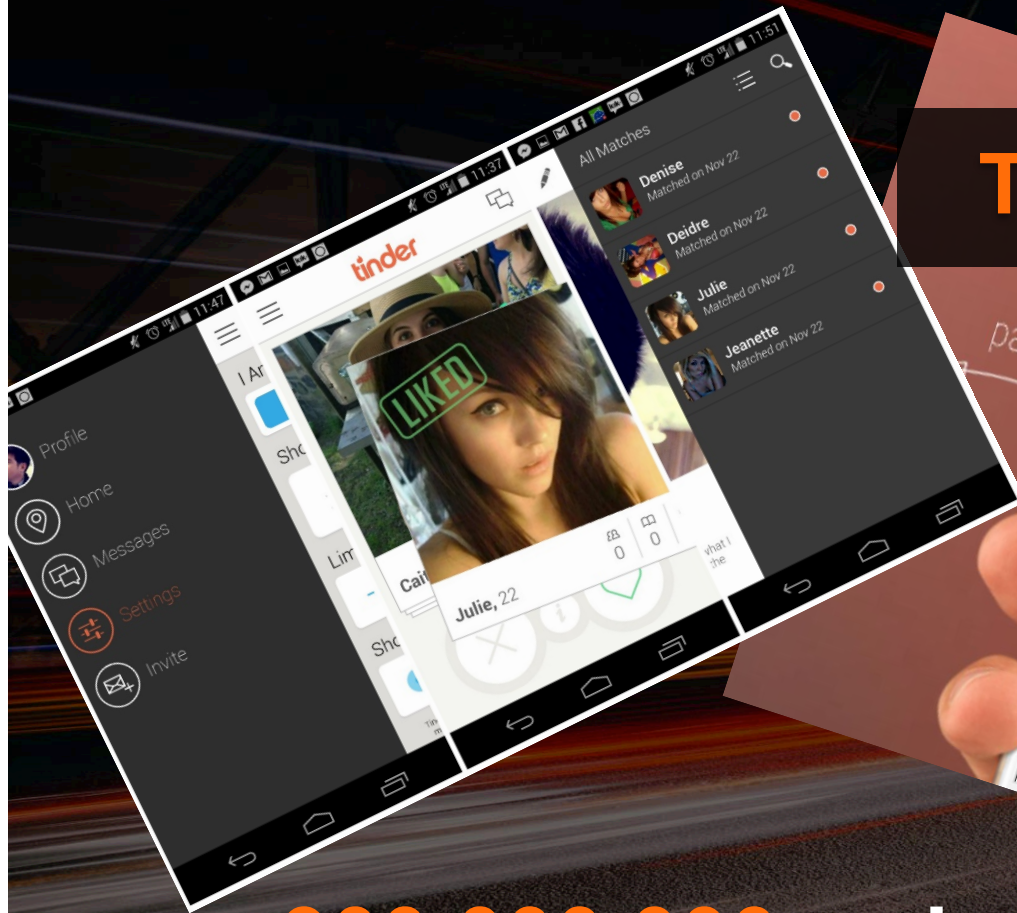
¥ 10,014,230,205

\$1,510,861,063 Total GMV in USD.
All GMV referenced is settled through Alipay.

无线成交占比 93%
GMV from Mobile



Transactions: Variety



900,000,000 swipes a day!



HyperScale Datacenter Powered by NAND Flash

- Higher QoS
- Lower TCO



WHY

HOW

WHAT

Open Channel SSDs for the datacenter

Applications

Middleware

Database

User

VFS

XFS, EXT4....

Block layer

NVMe Driver

Kernel

FTL

NAND Flash

Hardware



The Agent Problem

The Matrix, 1999

Open Channel SSD: **Reducing Agents**

Applications



- APIs tailored for applications
- Data hints sent to IO managers
- NAND status hints sent to Applications

IO and Flash
management

NAND

Open Channel SSD: **Eliminating Agents**

Applications with
IO and flash
management



Physical Block Addressing:
Channel/CE/Block/Page

NAND

OCS in **Kernel** Mode

User



Kernel

IO Engine (Block Device/
NS/KeyValue etc) and
flash manager

LightNVM Driver

Hardware

Controller logic (R/W/E)
+ NAND

OCS in **User** Mode

User



IO Engine (Block Device/
NS/KeyValue etc) and
flash manager

Kernel

LightNVM Driver

Hardware

Controller logic (R/W/E)
+ NAND



OCS Become Integrated Part of Applications



WHY



HOW

WHAT

Open Channel SSDs for the datacenter

Introducing Industry's First Commercialized OCS Platform

The "Venice"



Partnership on OCS

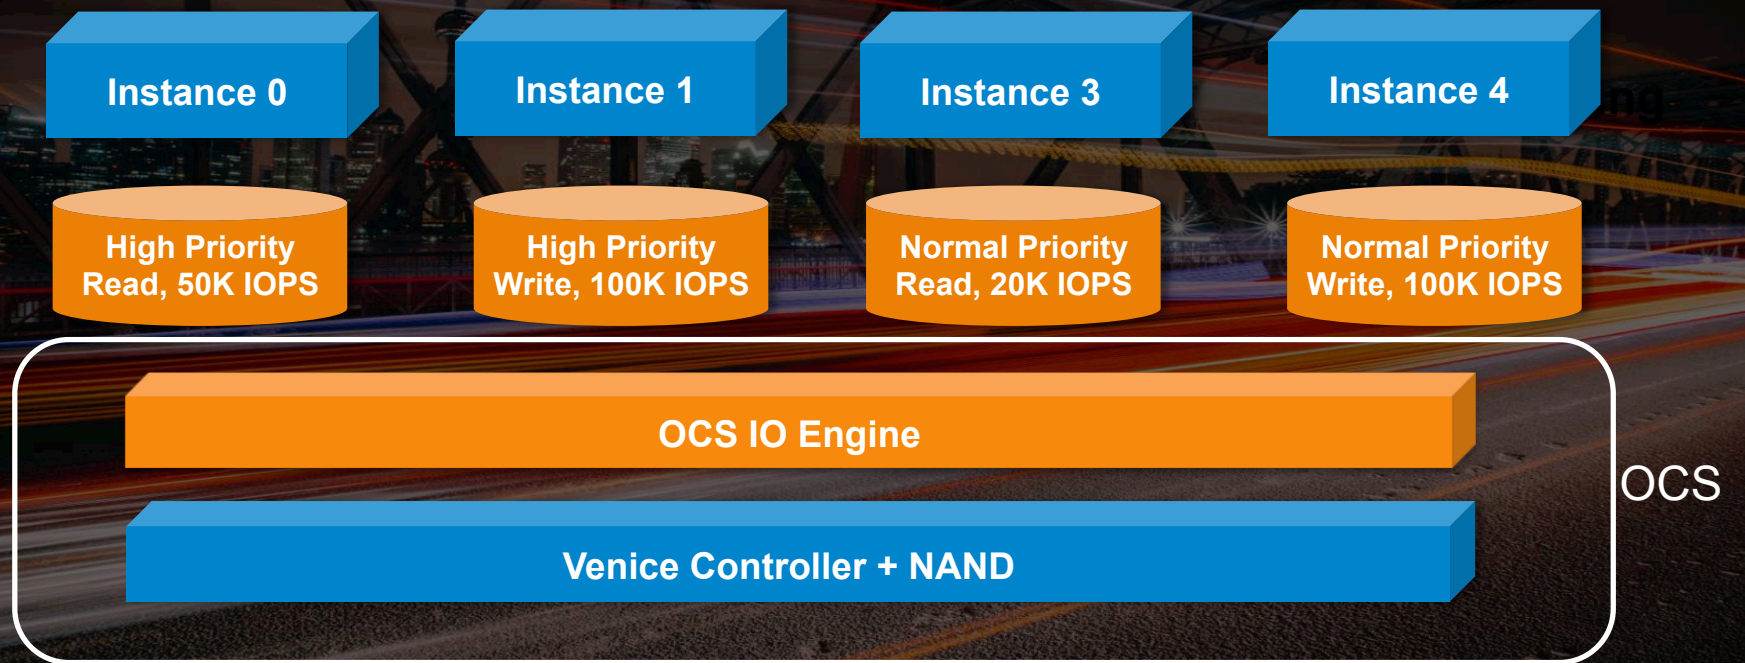
conforming to the Alibaba Open Channel SSD Spec



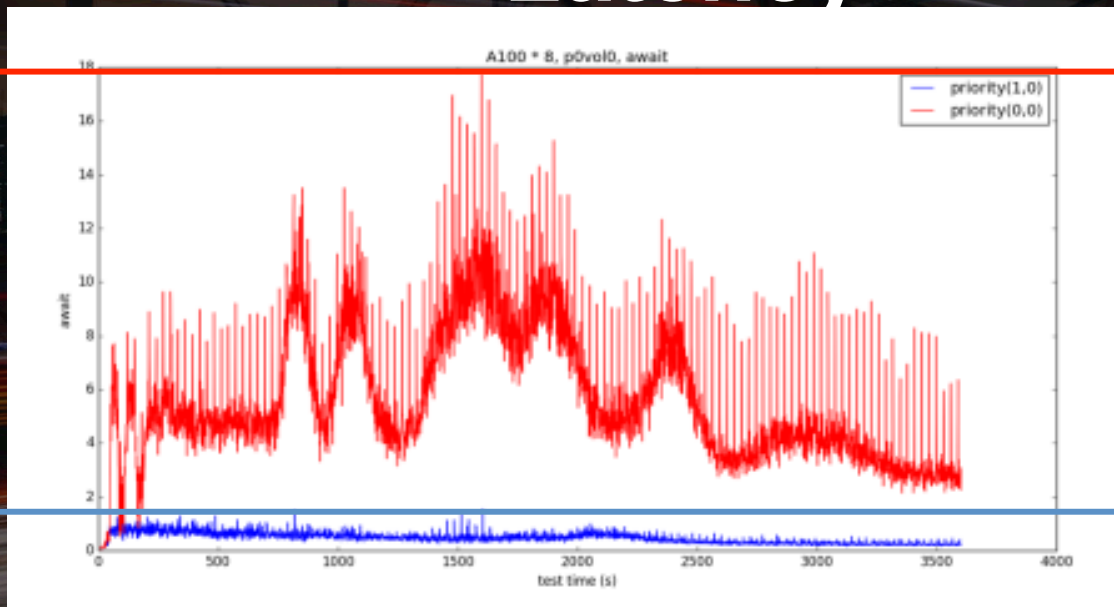
Shannon Systems

In production today

OCS for the Cloud: QoS Provisioning

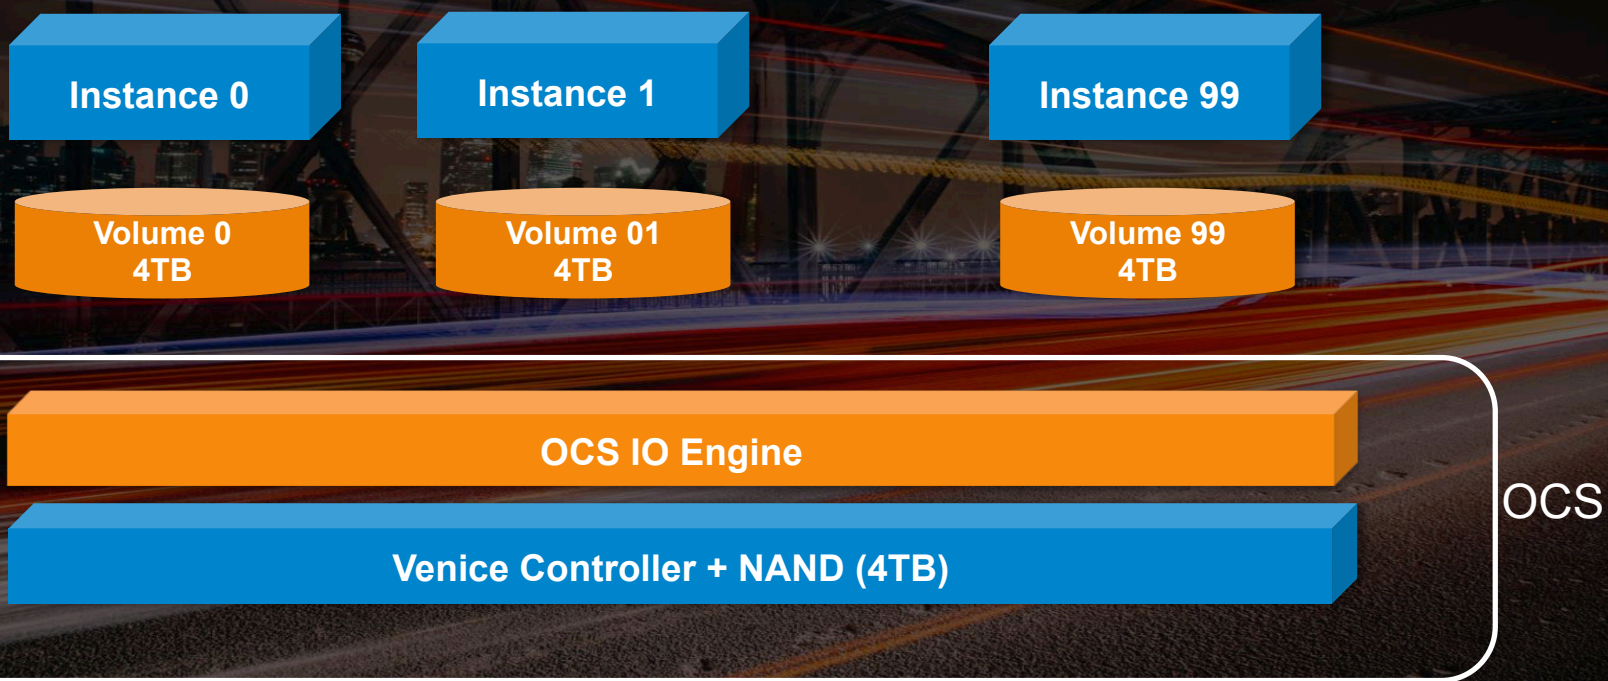


Prioritized IO Reduces Latency

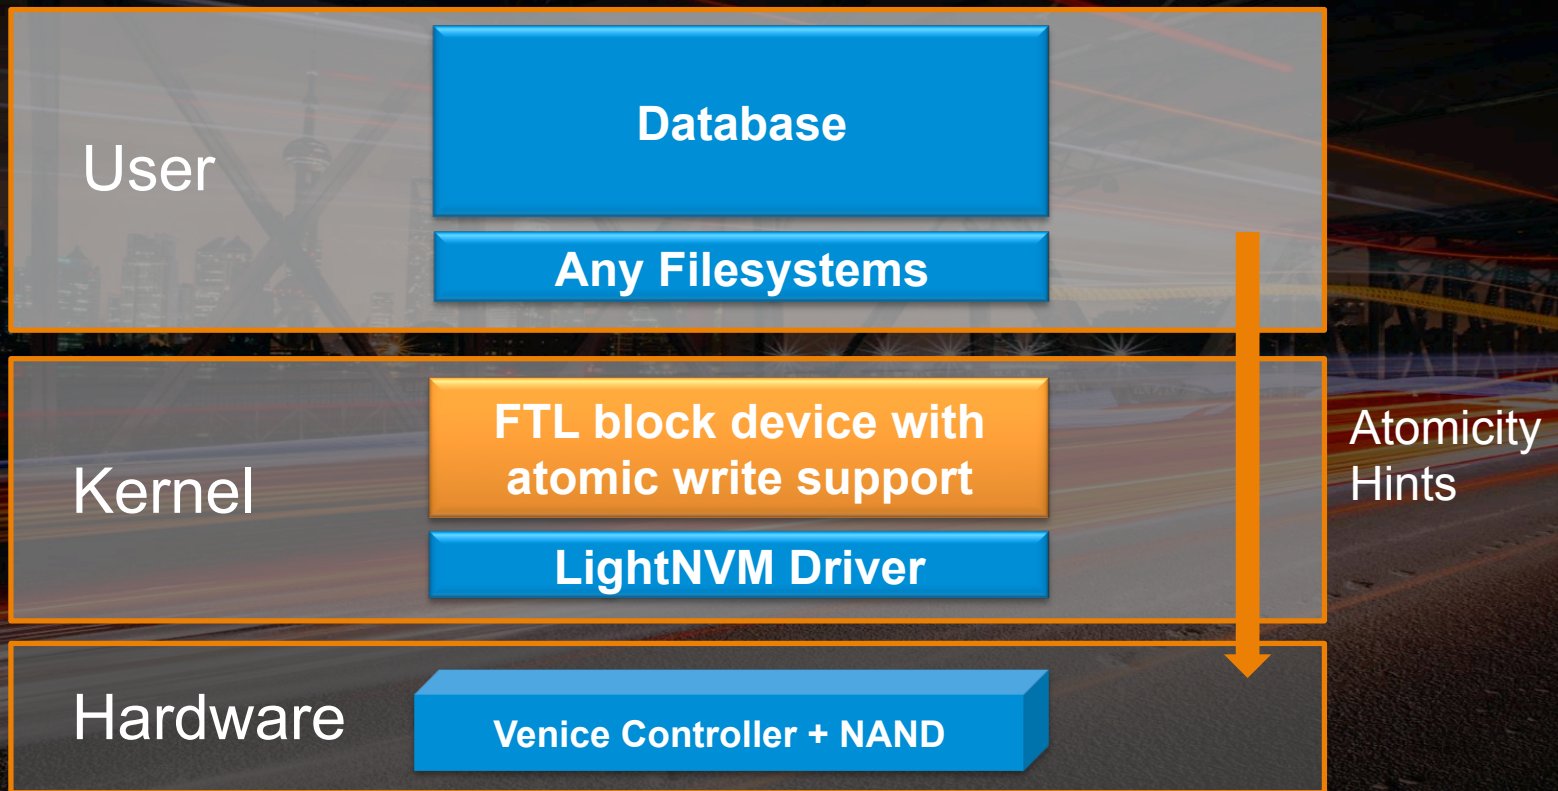


90% reduction

OCS for the Cloud: **Thin Provisioning**



OCS for Databases: Atomic Writes



OCS for Databases

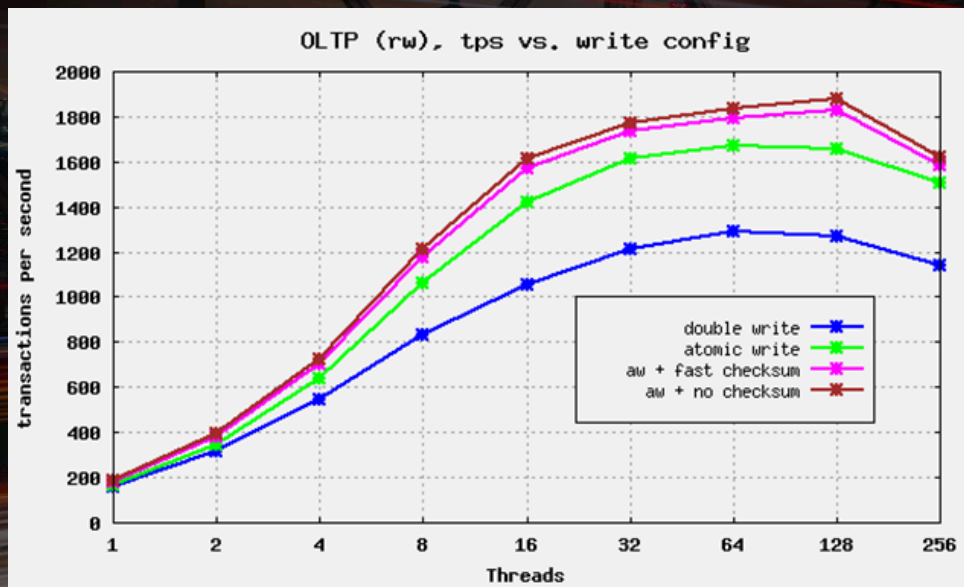
Atomic Writes



Shannon Systems

Native support after Ver 10.2

OCS Atomic Writes – MariaDB® Support



- TPS increases by up to **50%**

<https://mariadb.com/kb/en/library/changes-improvements-in-mariadb-102/>

OCS for NoSQL: **Key-Value**

- KV access
- Kernel/user-space mode
- Physical based addressing

KV Applications

User-space KV Engine

Venice Controller + NAND

OCS for NoSQL: **Key-Value**



KV-SSD: Breakthrough Performance



Random 8KB value size access

Partnership on KV-SSD



Shannon Systems

Key Takeaways

- **OCS have arrived**
- OCS enable applications to **directly manage flash**, bypassing agents and OCS become part of applications
- OCS offer significant opportunity for **QoS improvement** for **datacenter** applications:
 - Cloud and virtualization
 - Databases
 - Key value store

Come to Visit US at Booth #219

