Data Explosion Demands Ethernet Fast Flash Fabrics

Ahmet Houssein
VP of Marketing & Strategic Development
Solarflare

Data Centers in the Past

Roadways
Networks for enterprise workloads

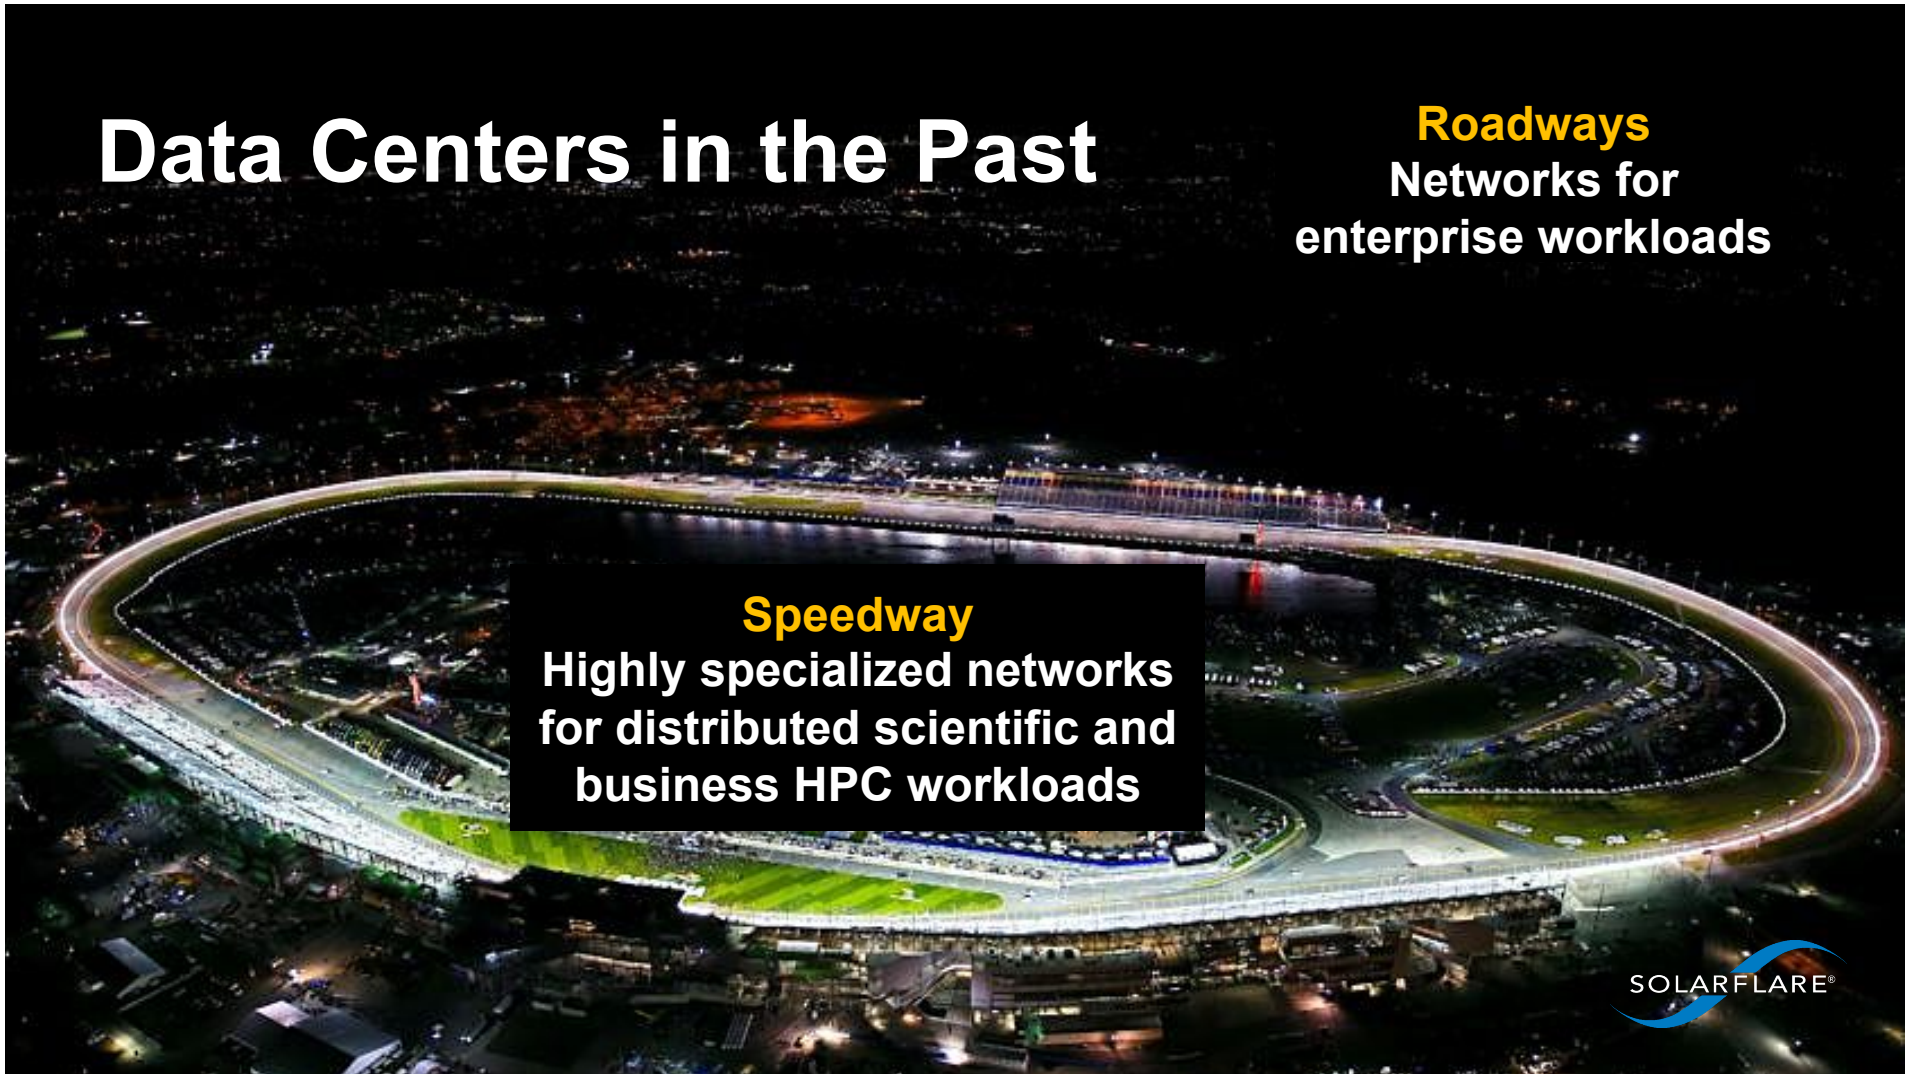

Speedway
Highly specialized networks for distributed scientific and business HPC workloads

SOLARFLARE®

Modern Data Centers

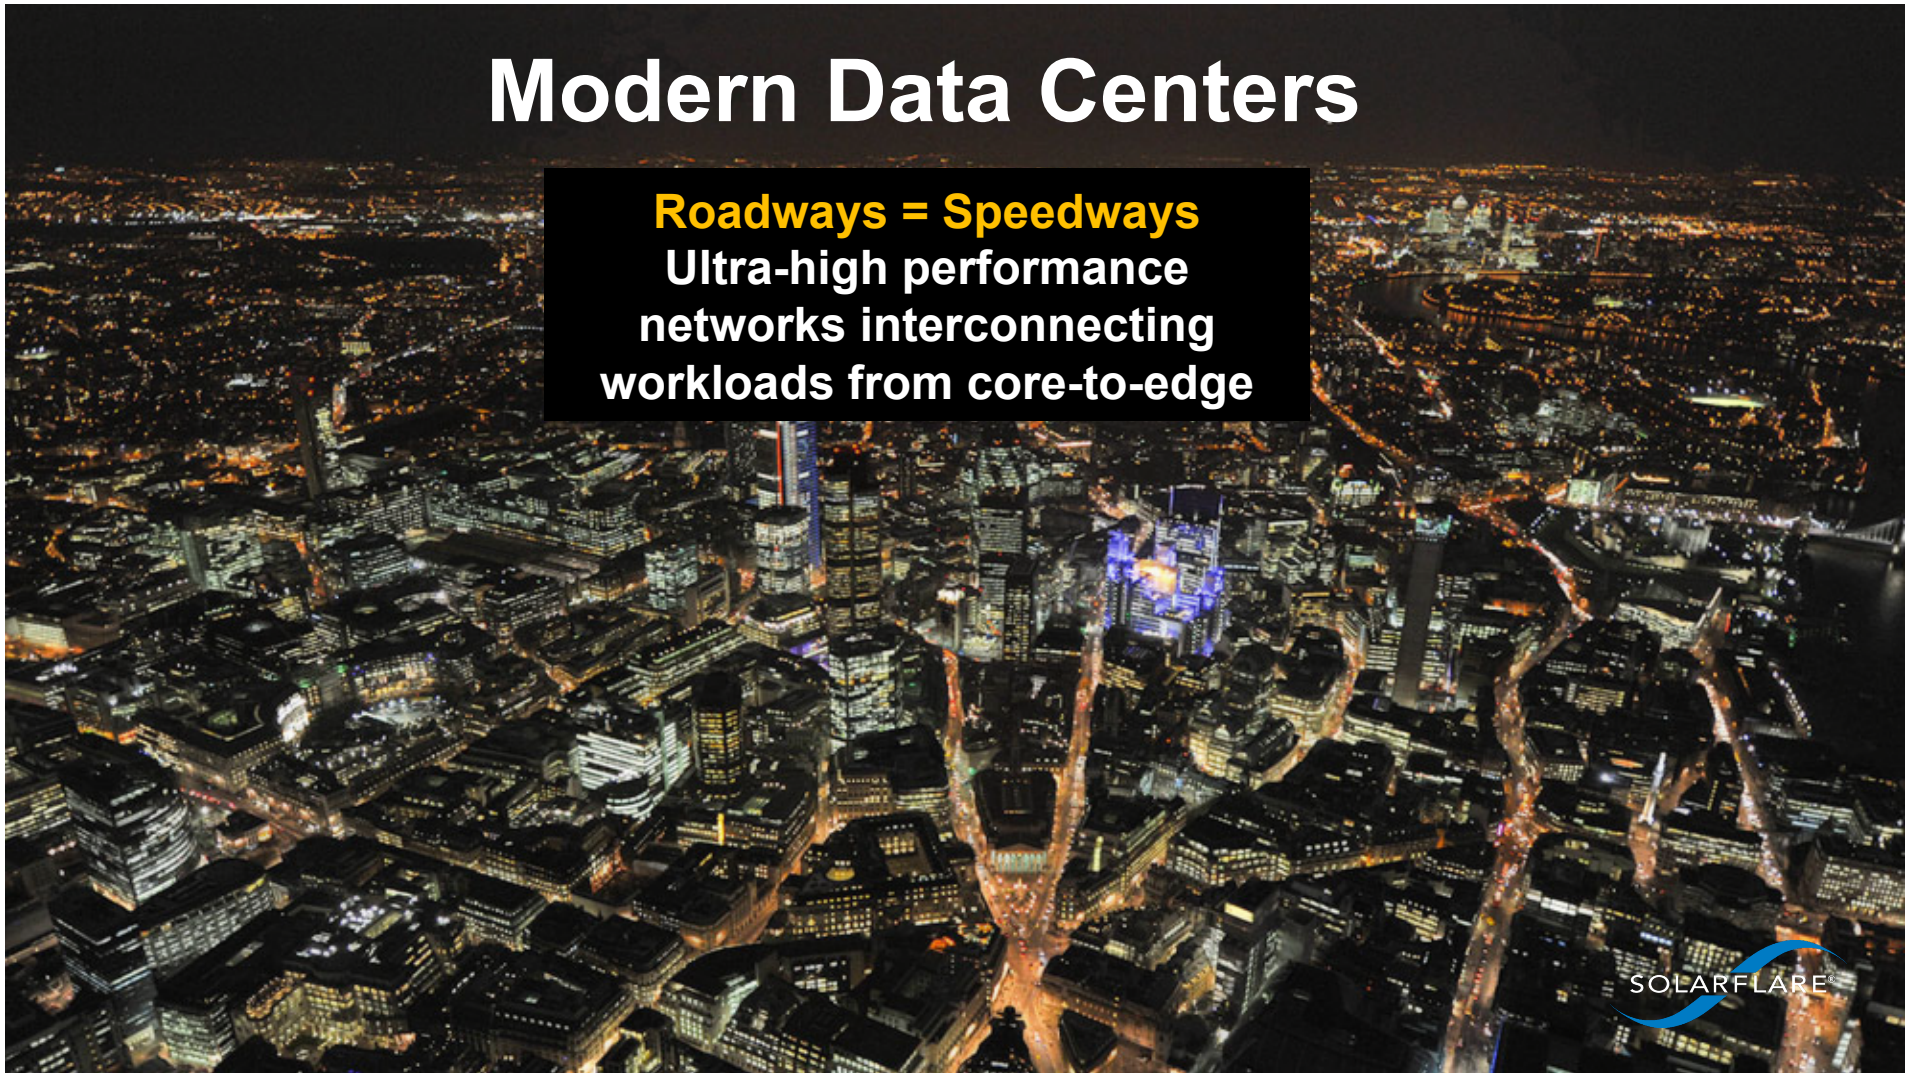

NVMe Flash Storage in the Data Center

SSD: Sparks explosive growth. SAS/SATA Storage bottleneck for servers.

NVMe: Addresses interface bottleneck for servers. Limited to inside the servers.

NVMe-oF: FC, IB & RoCE limit adoption and add cost.

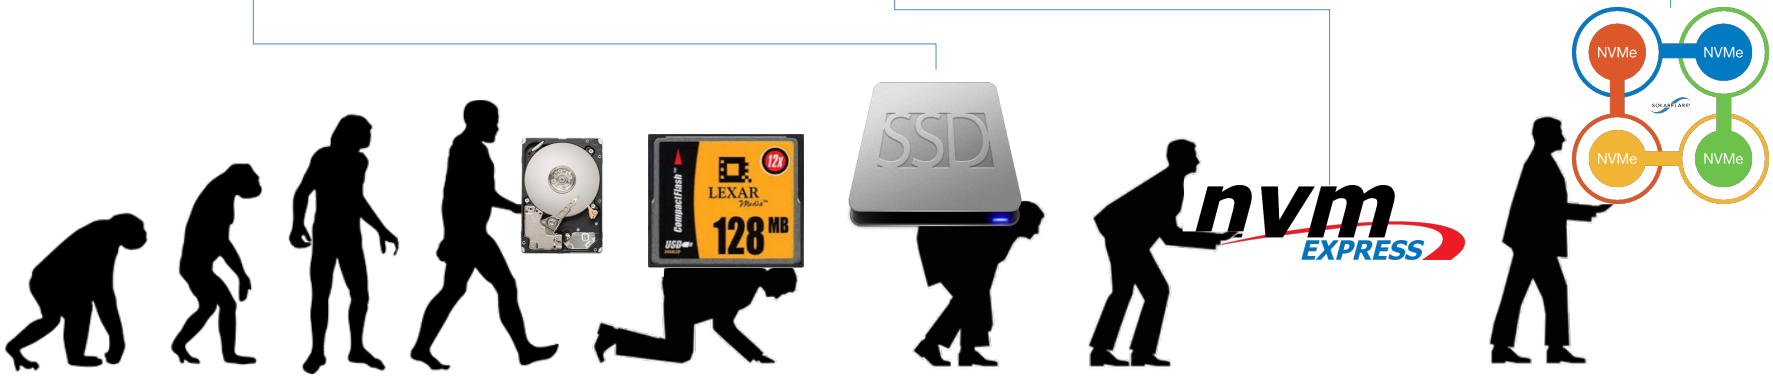

Solarflare Mission

Empower all applications to gain benefits of FLASH storage without cost and implementation issues by enabling NVMe flash with ubiquitous TCP/IP

SOLARFLARE®

Empirical Data on NVMe over TCP IOPS

SOLARFLARE®

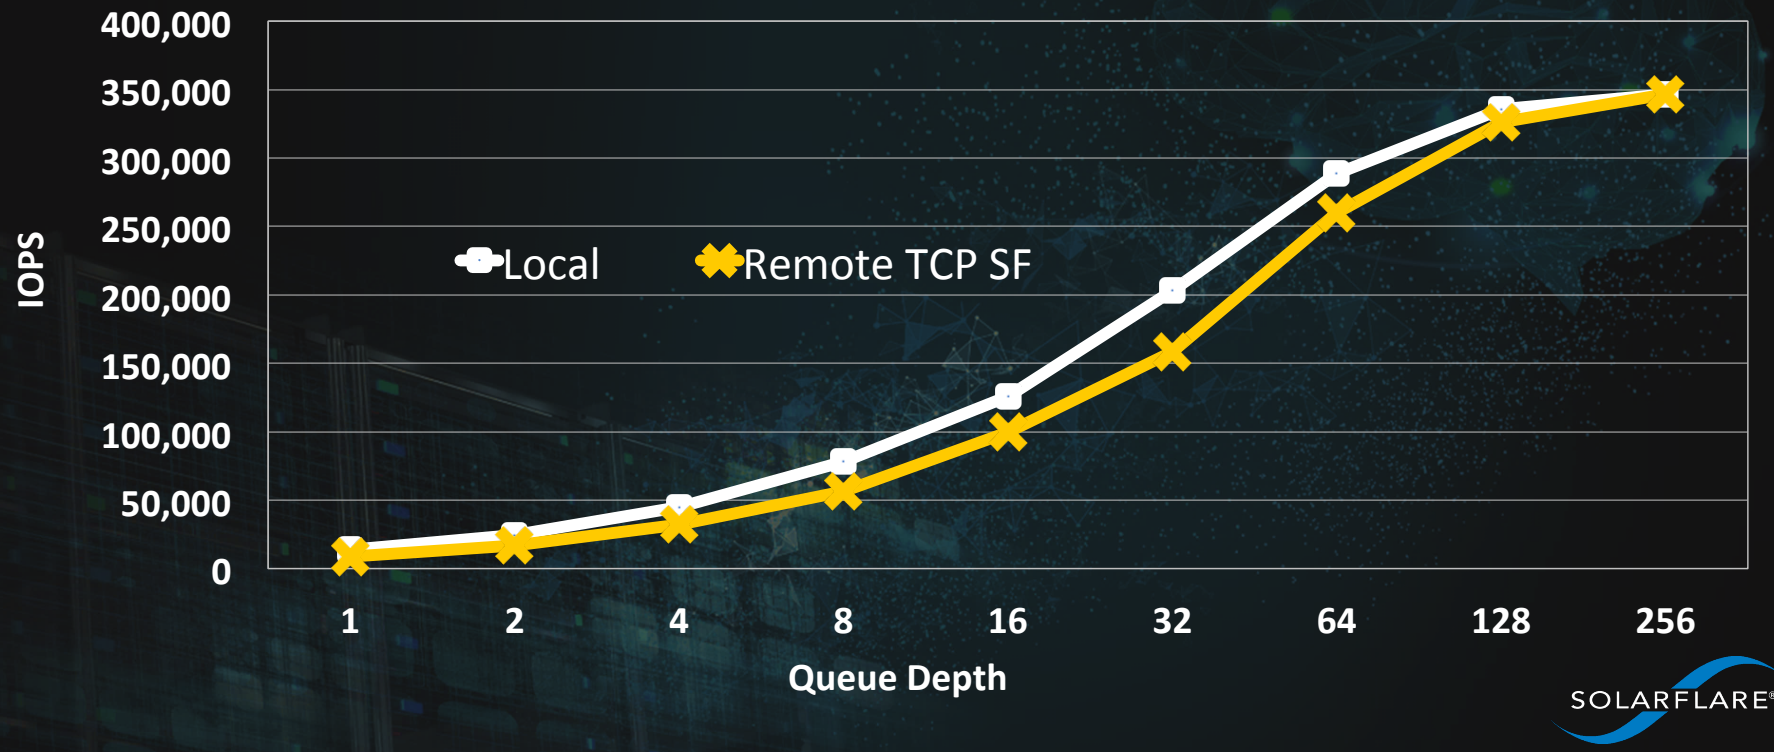

SOLARFLARE®

Empirical Data on NVMe over TCP Latency

SOLARFLARE®

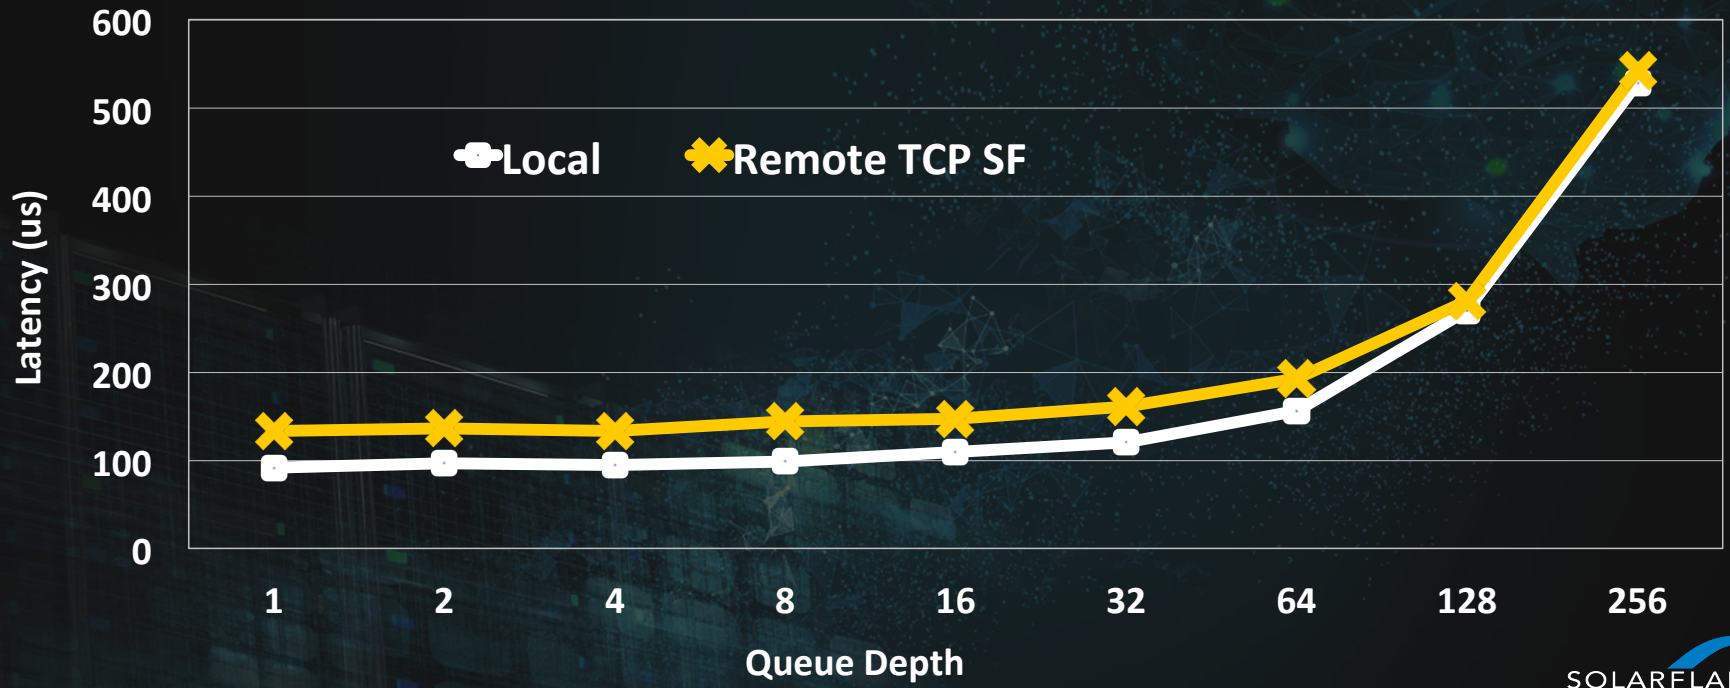

SOLARFLARE®

Empirical Data on NVMe over TCP Bandwidth

SOLARFLARE®

Sustained Random Read IO, Queue Depth = 256

(Kernel NVMe over TCP)

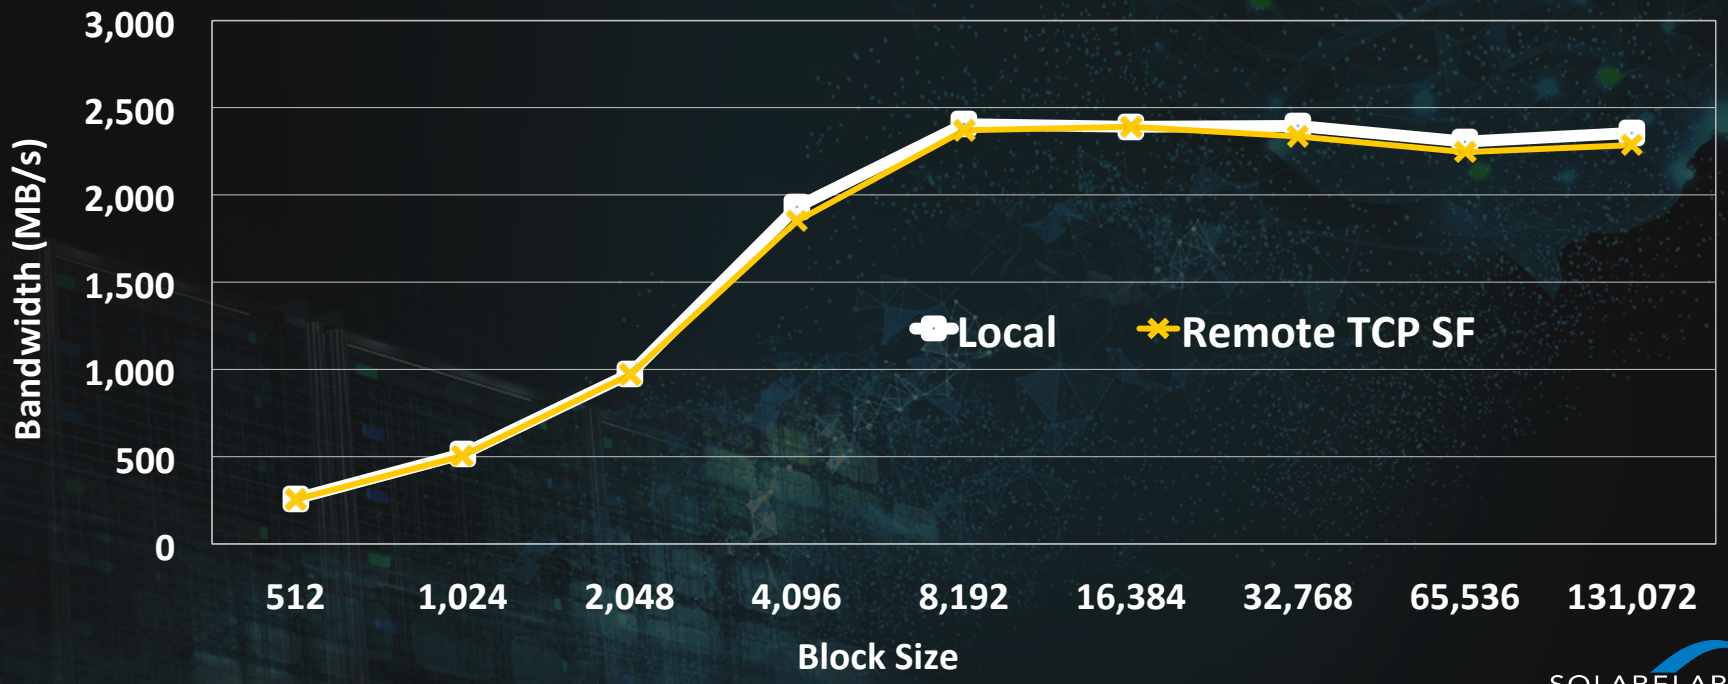

SOLARFLARE®

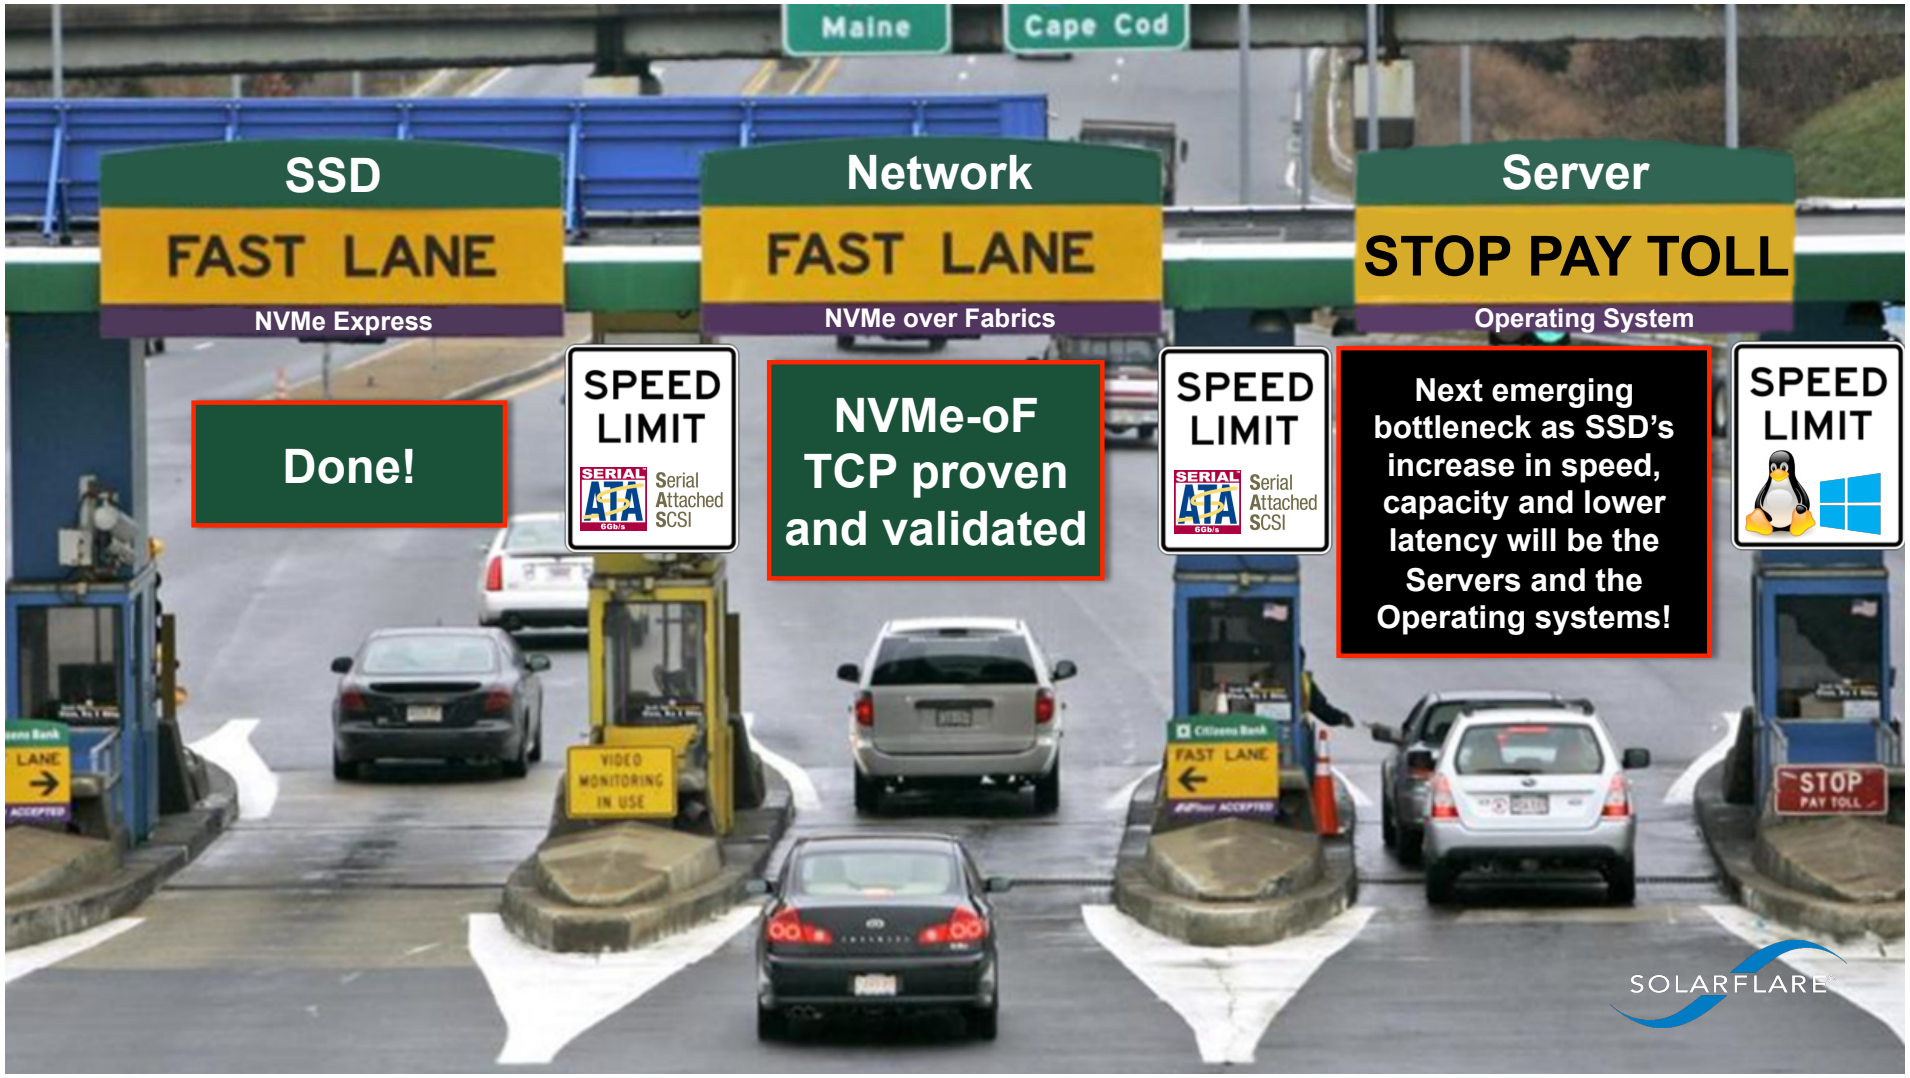
Maine
Cape Cod
SSD
FAST LANE
NVMe Express
Done!
SPEED LIMIT
Serial Attached SCSI
Network
FAST LANE
NVMe over Fabrics
NVMe-oF TCP proven and validated
SPEED LIMIT
Serial Attached SCSI
Server
STOP PAY TOLL
Operating System
Next emerging bottleneck as SSD's increase in speed, capacity and lower latency will be the Servers and the Operating systems!
SPEED LIMIT
VIDEO MONITORING IN USE
FAST LANE
STOP PAY TOLL
SOLARFLARE

IoT Dynamite

Filling Big Data Stores & Waiting for Edge Analytics

212
Billion
Sensor enabled objects by 2020 - IBM

30
Billion
Sensor enabled objects networked by 2020 - IBM

SOLARFLARE®

A 5G Match

The Key to Real Time Analytics in Micro Data Centers

1980s	1990s	2000s	2010s	2020s
1G	2G	3G	4G	5G
Analog Voice	Digital Voice	Mobile Broadband	Faster Better	Computing Platform

10,000 x more traffic • 100x more devices • Less than 1 ms latency

10 year battery life • 10 Gb/second data rate

SOLARFLARE®

An Explosion in Real Time Analytics

Traditional Data Warehouse

Structured Data
Long-Term Analytics
Planning to Make Decisions

Big Data Store

Unstructured Data
Short-Term Analytics
Make a Decision Today

Edge Devices

Streaming Data
Real-Time Analytics
Make a Decision **NOW**

SOLARFLARE®

Servers in 2018 capacity Dramatic Increase

100 Servers x 72 CPU Cores = 7,200 Interconnected Cores

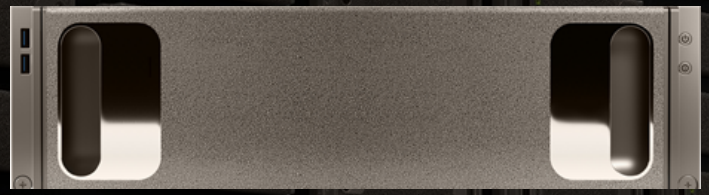

CPUs	1x Xeon Phi
Cores	72

SOLARFLARE®

A New Era of Networked Storage

With new class of SSD's like Intel Optane reducing latency to~10 us will have a significant impact!

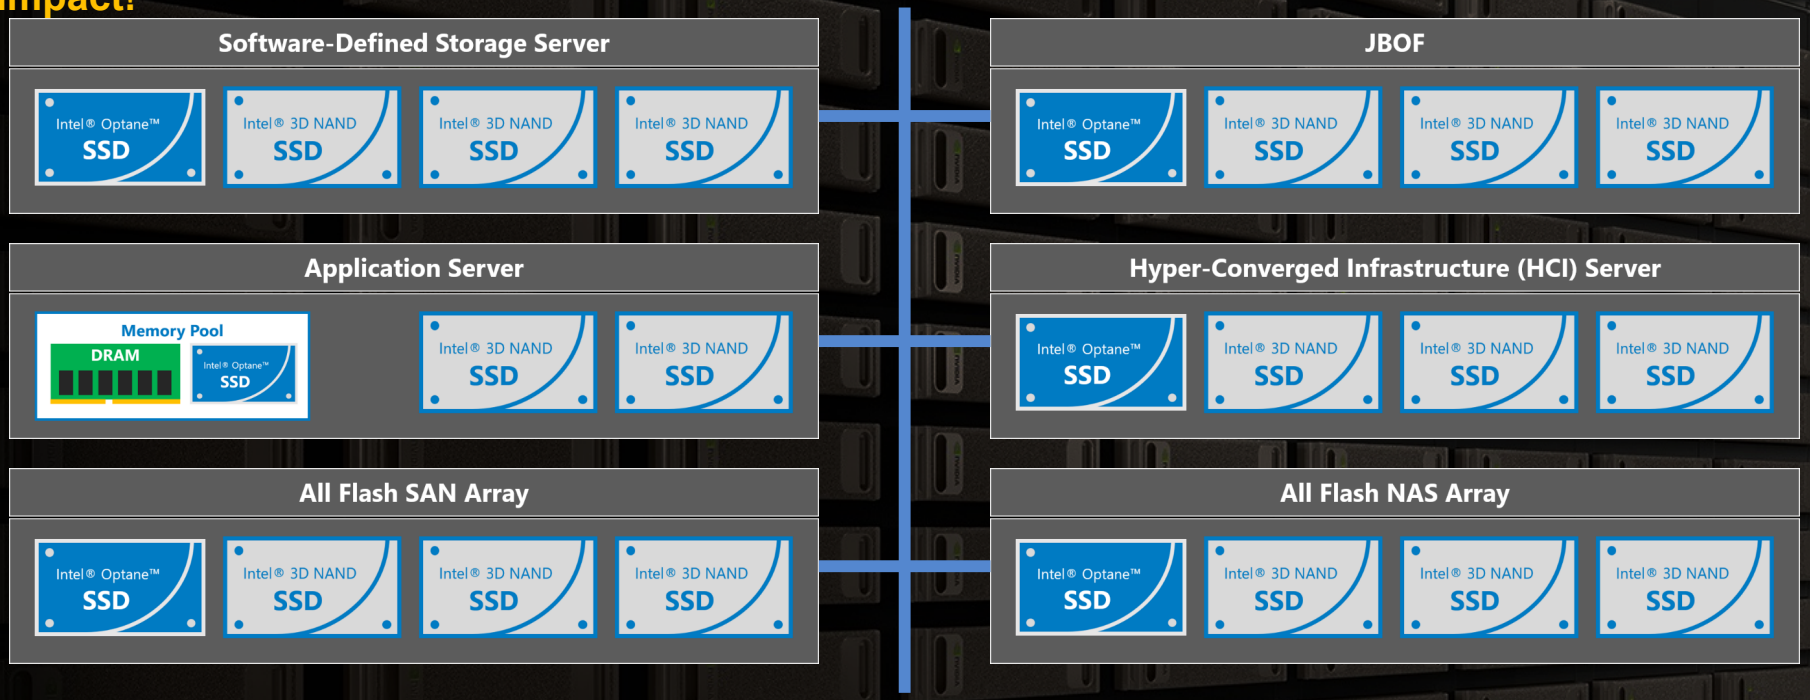

Microsecond Access Anywhere in the Data Center with Standard Ethernet

2018 Flash Storage Bottlenecks

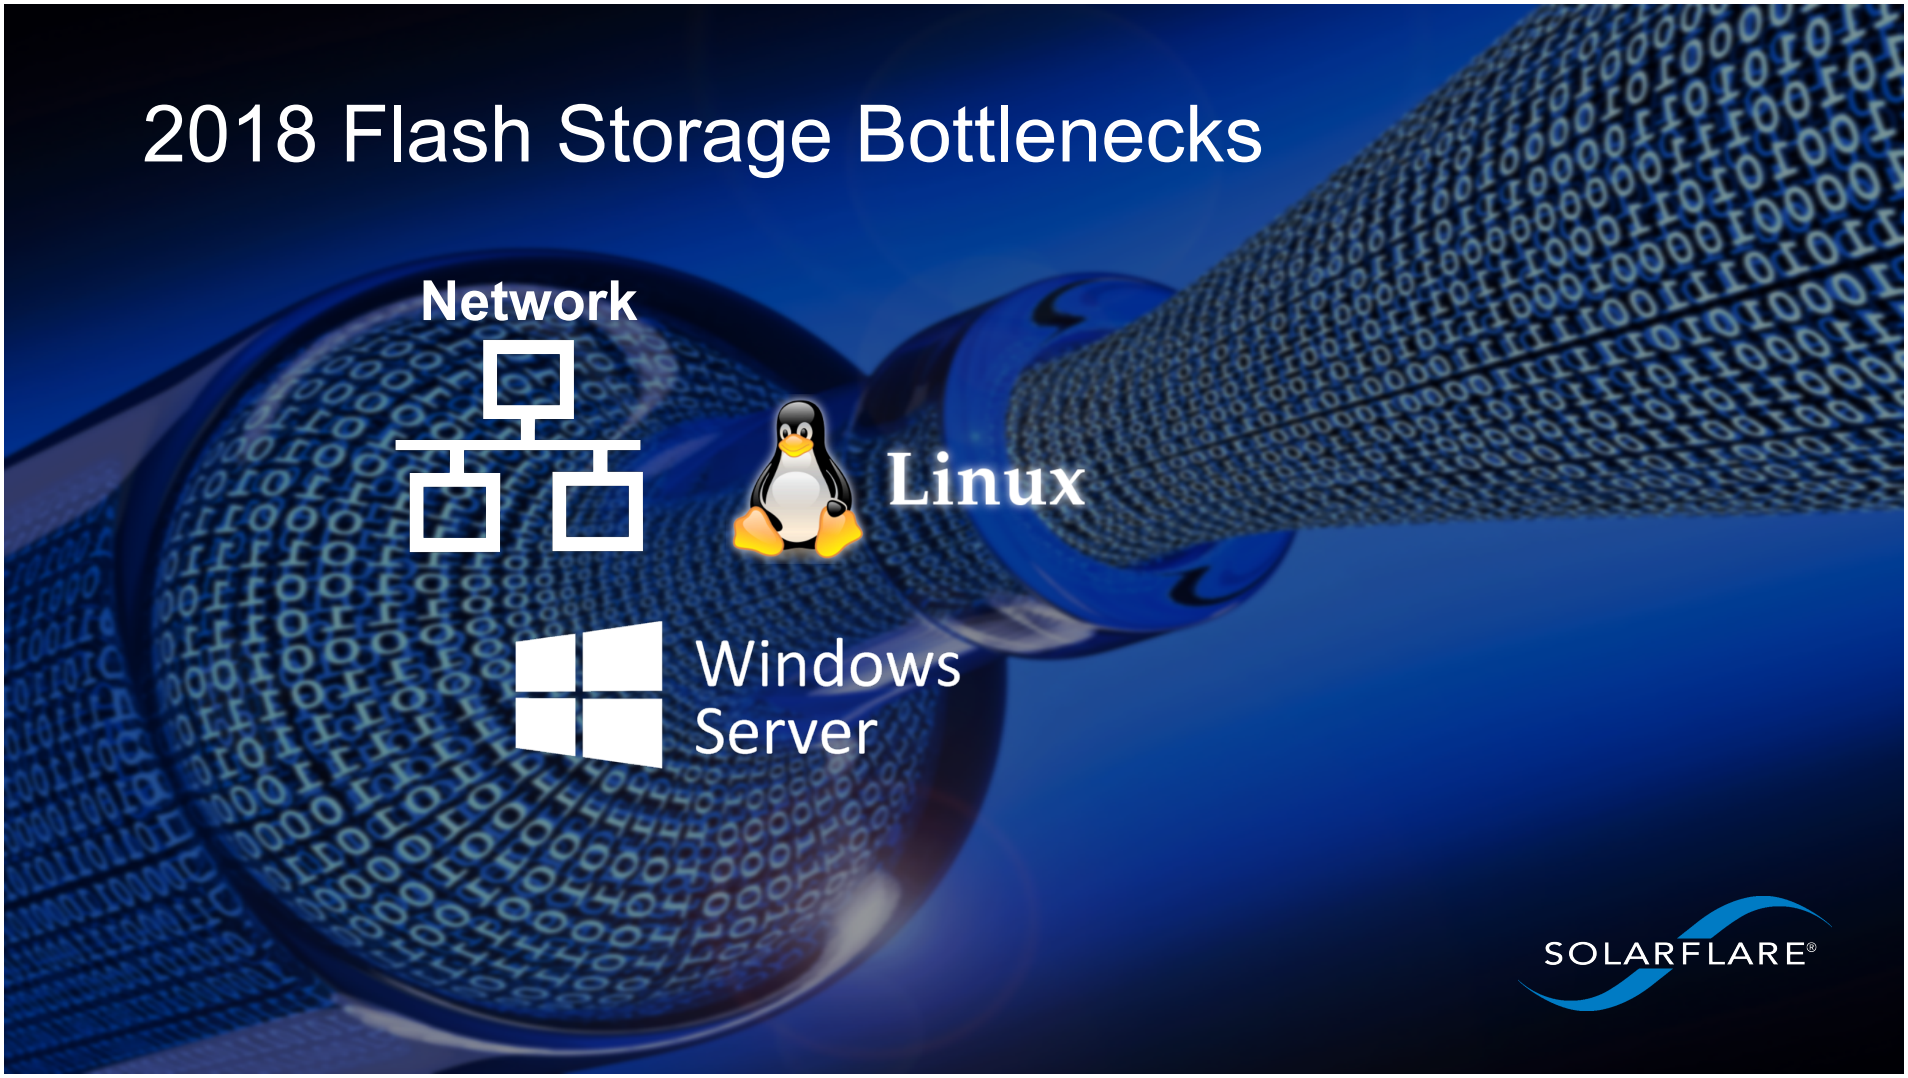

Under the Hood of NVMe over TCP

Inside The Server Bottleneck

SOLARFLARE®

The Problem: OS Context Switching

Common problem that occurs when many micro services are sharing storage devices is operating system context switching, buffer copying, the constant suspending and resuming of processes which kills application performance.

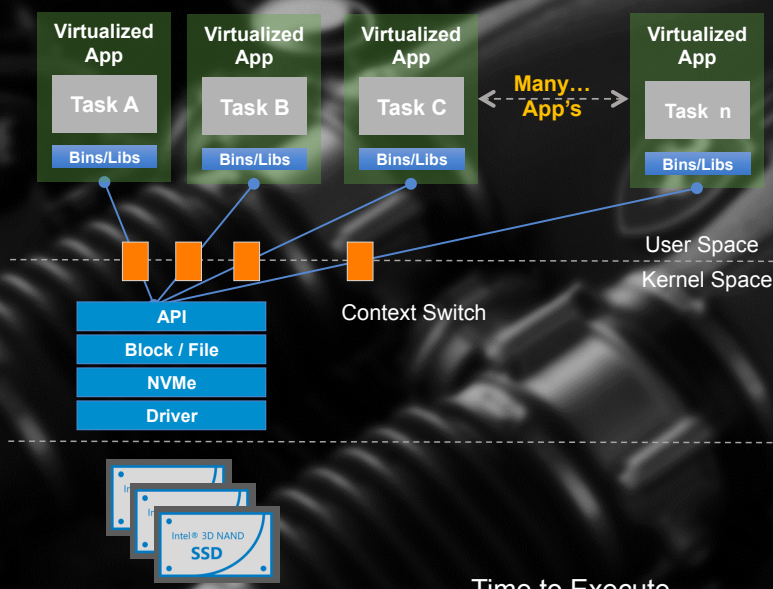

The Solution: User Space Block Storage (SPDK)

SPDK moves Block/NVMe layer in to user space eliminating context switching, buffer copying and blocking. Significantly removing overhead reducing latency and increasing scalability.

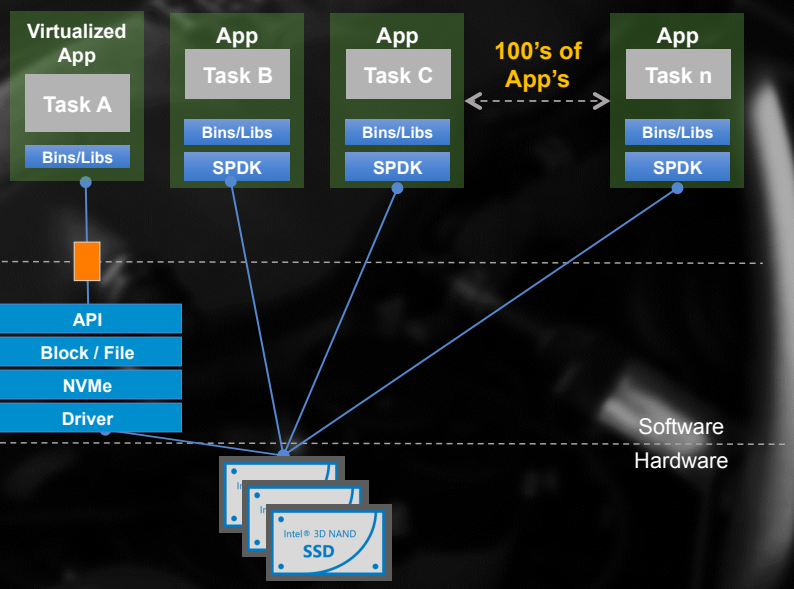

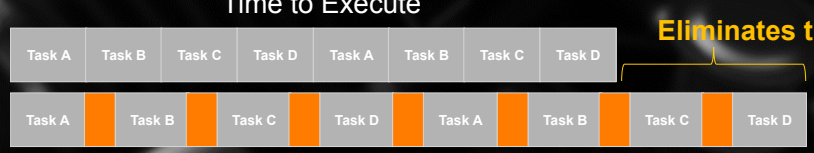

The Networking Bottleneck

SOLARFLARE®

The Solution: User space Block Interface (SPDK) and TCP/IP (Onload) Stack

By bypassing Operating System for both NVMe and NCMe-oF TCP eliminates Bottleneck

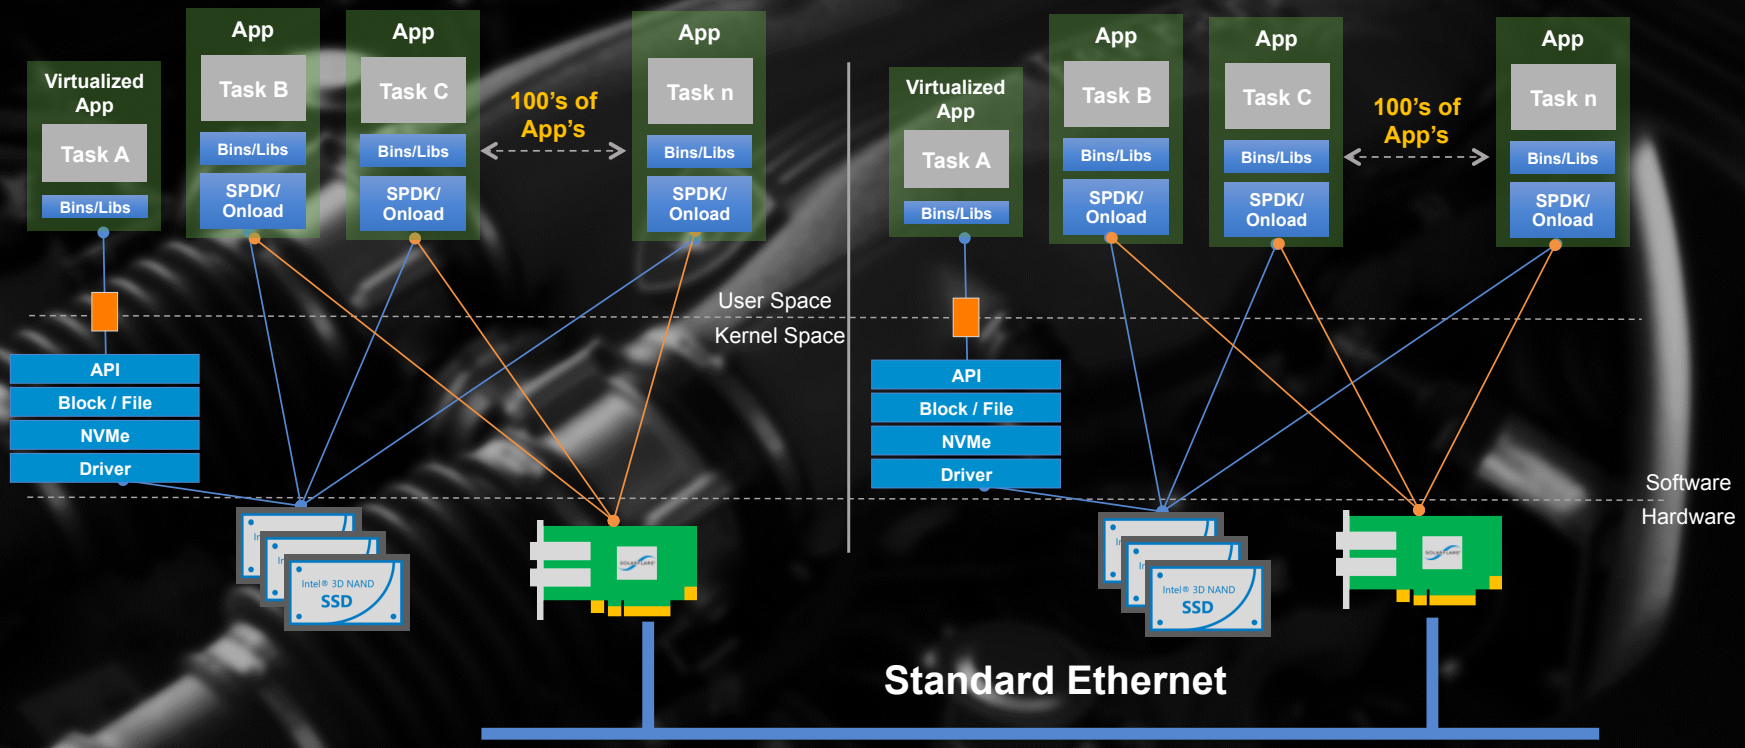

User Space Solutions Drive Performance & Scale

Intel Optane SSD and NVMe over TCP

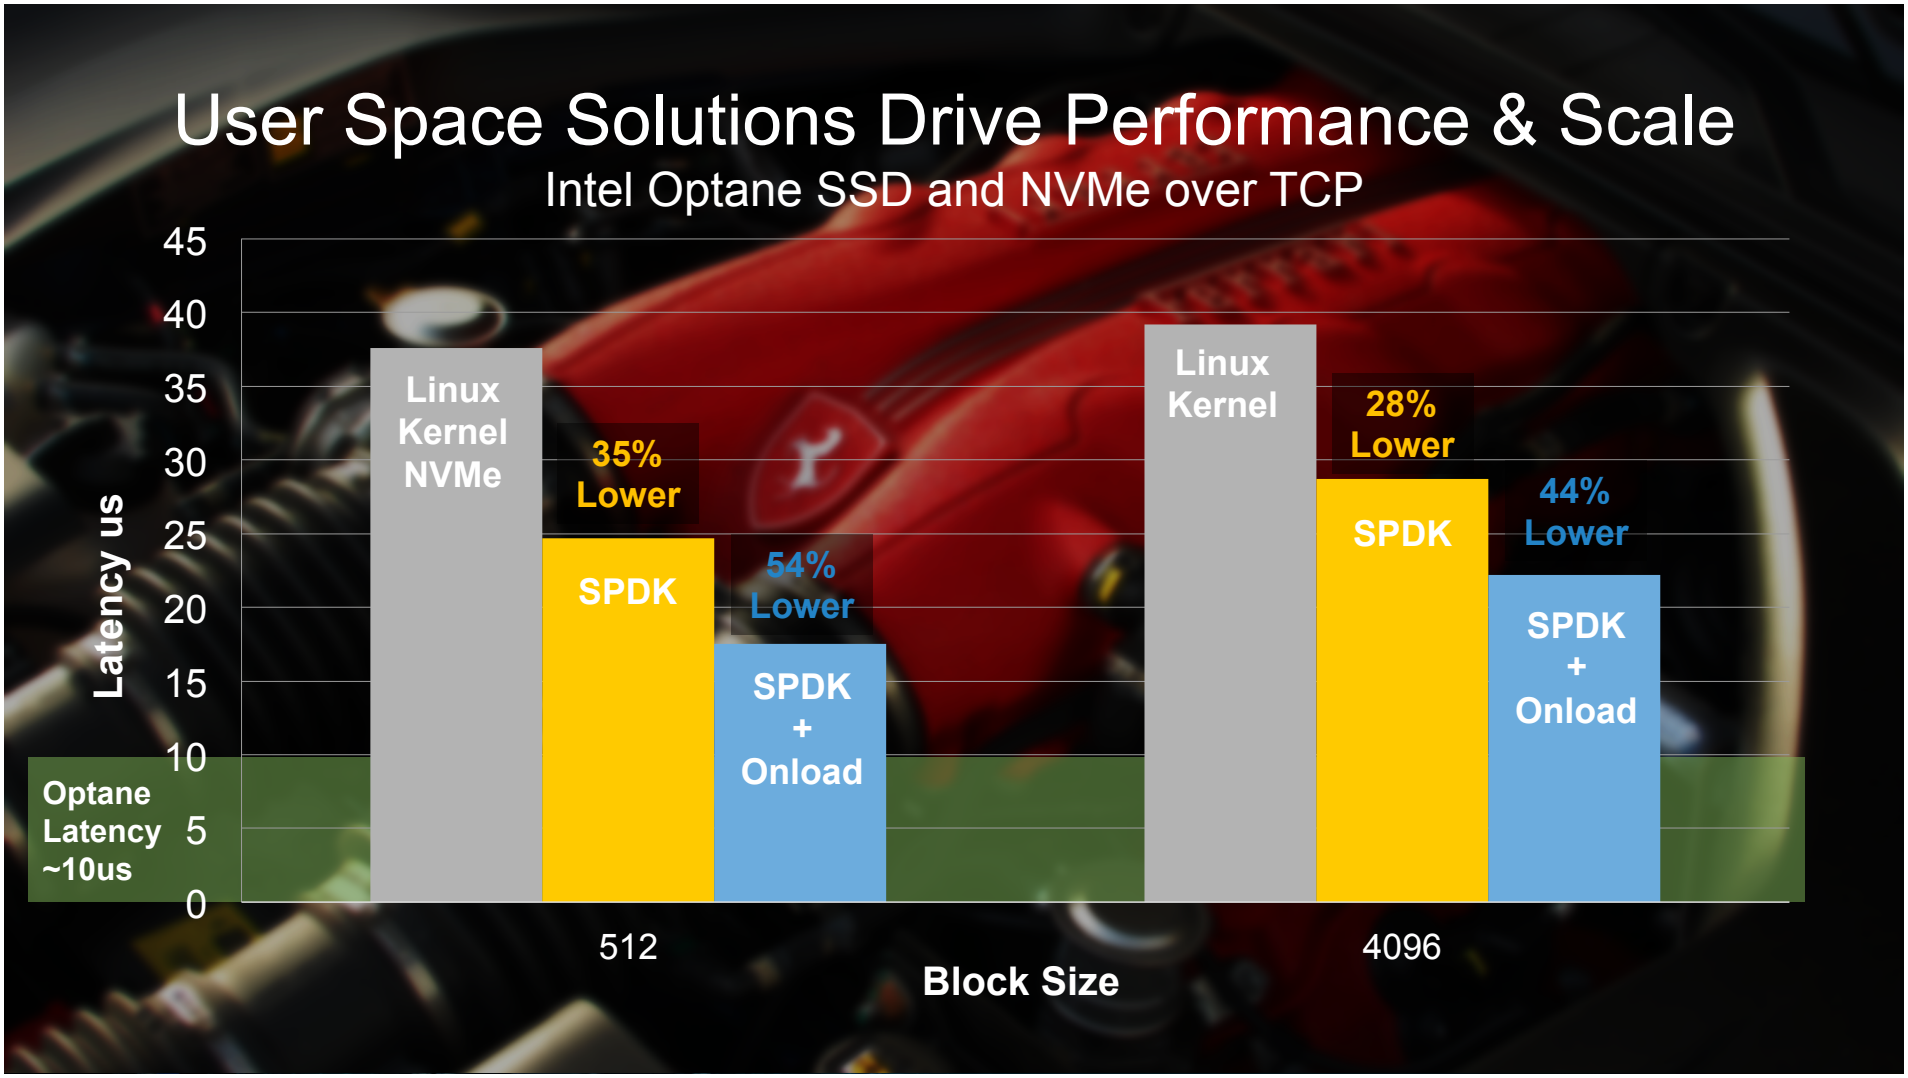

NVMe over TCP: Delivers on Promise

Simple to Deploy and Use on Existing Network

98% of installed servers connected with TCP

2% of installed servers connected with RDMA

SOLARFLARE®

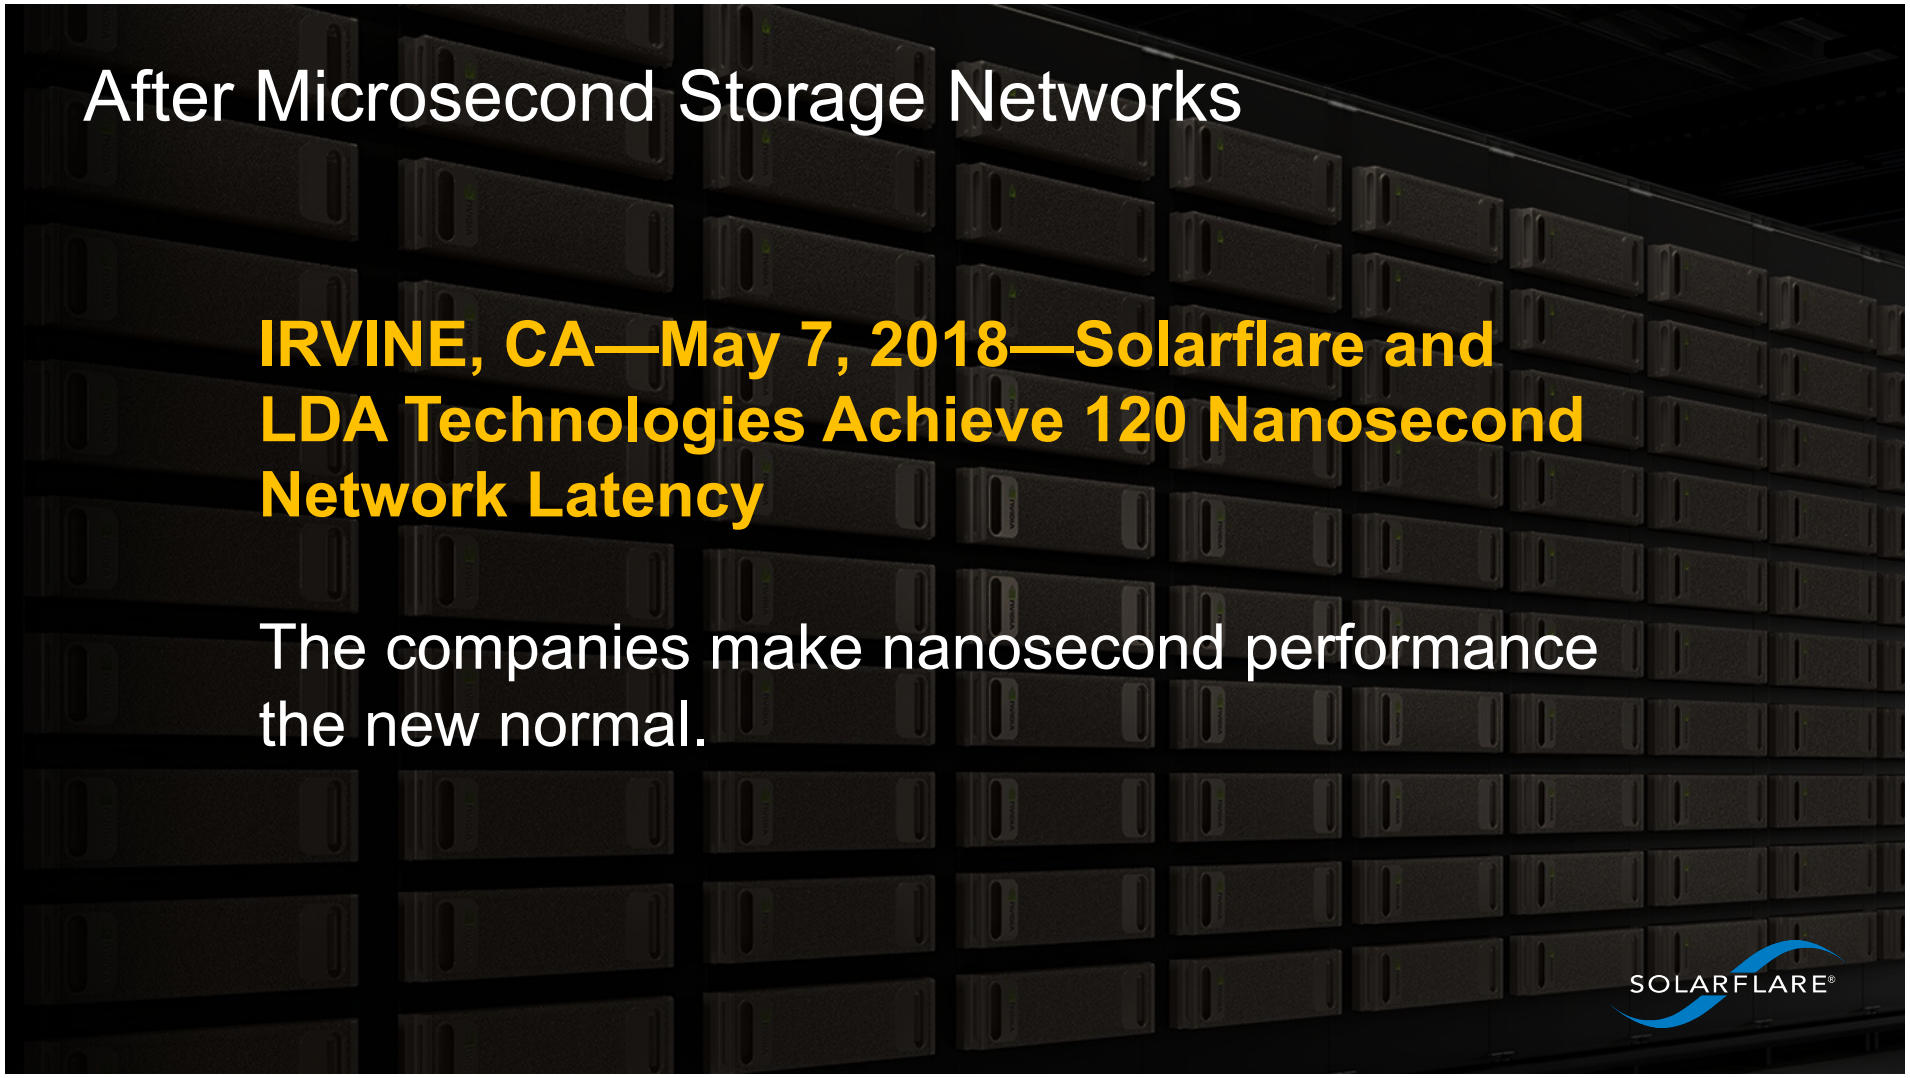
After Microsecond Storage Networks
IRVINE, CA—May 7, 2018—Solarflare and LDA Technologies Achieve 120 Nanosecond Network Latency
The companies make nanosecond performance the new normal.
SOLARFLARE®

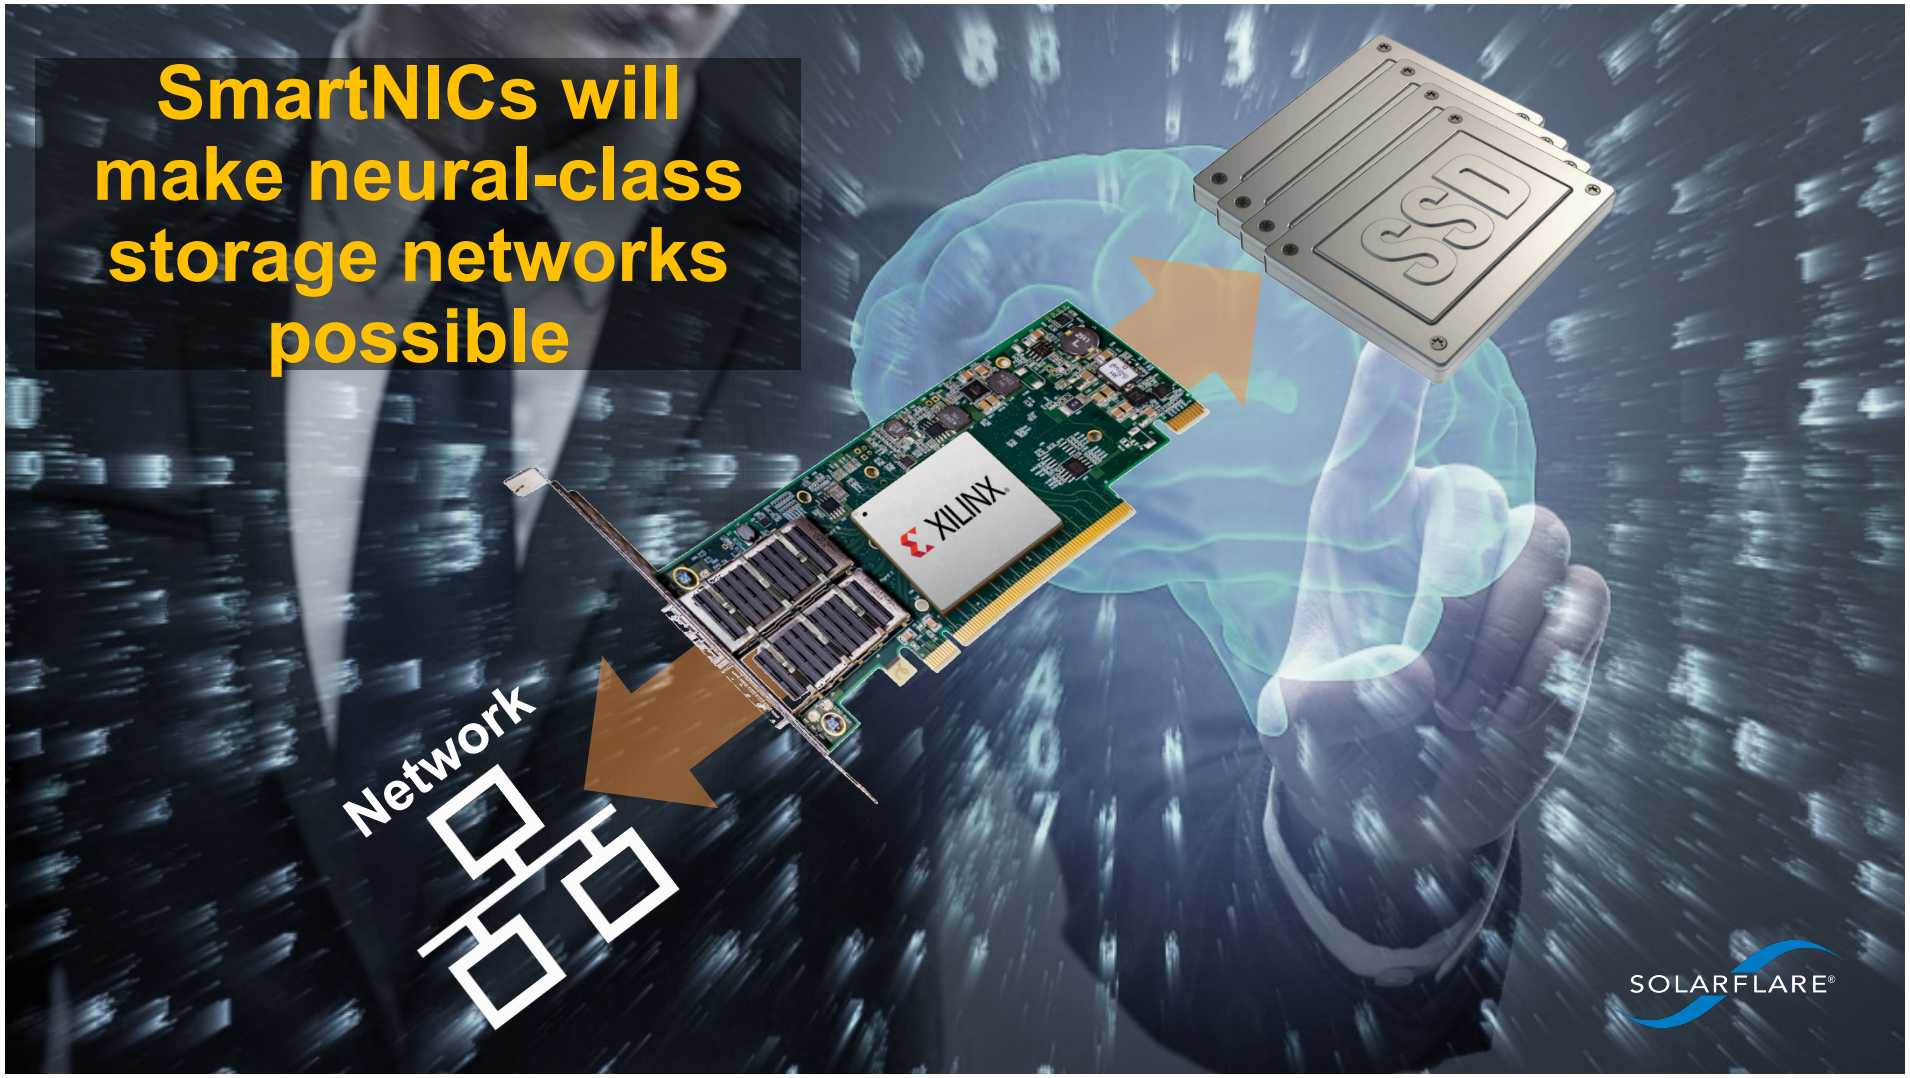
SmartNICs will make neural-class storage networks possible
SSD
XILINX
Network
SOLARFLARE®


SOLARFLARE®
Neural-Class Networks Made Simple