SDC 19

September 23-26, 2019
Santa Clara, CA

Security, Integrity and Choices for NVMe over Fabrics

Nishant Lodha
Marvell

Agenda

- NVMe-oF®, the choices and the confusion
- Use Cases by Fabric
- Securing NVMe-oF
- Key Takeaways

NVMe-oF

Scaling our NVMe Requires a (Real) Network

- Many options, plenty of confusion
- Fibre Channel is the transport for the vast majority of today's all flash arrays
    - FC-NVMe Standardized in Mid-2017
- RoCEv2, iWARP and InfiniBand are RDMA-based but not compatible with each other
    - NVMe-oF RDMA Standardized in 2016
- FCoE fabric is an option
- NVMe/TCP – is here! Standardized in NOV2018

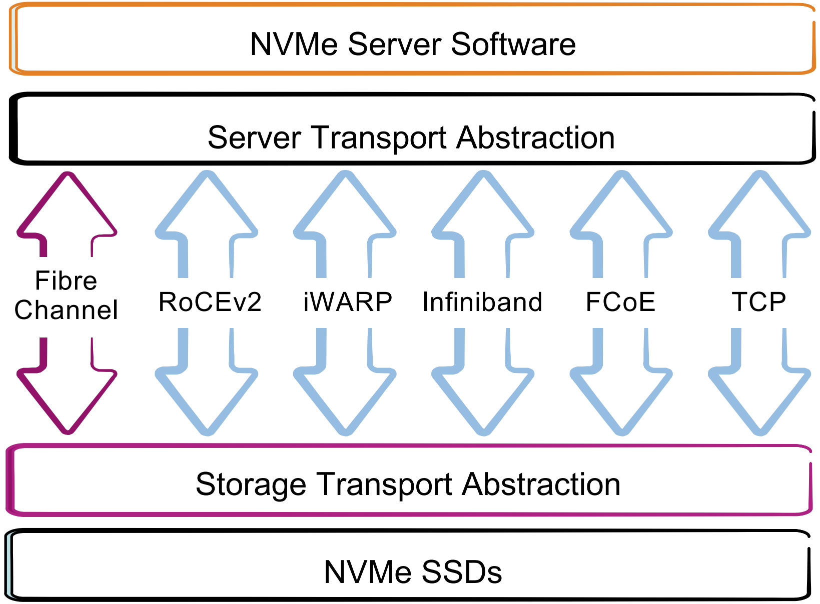

RDMA Use Cases by Application

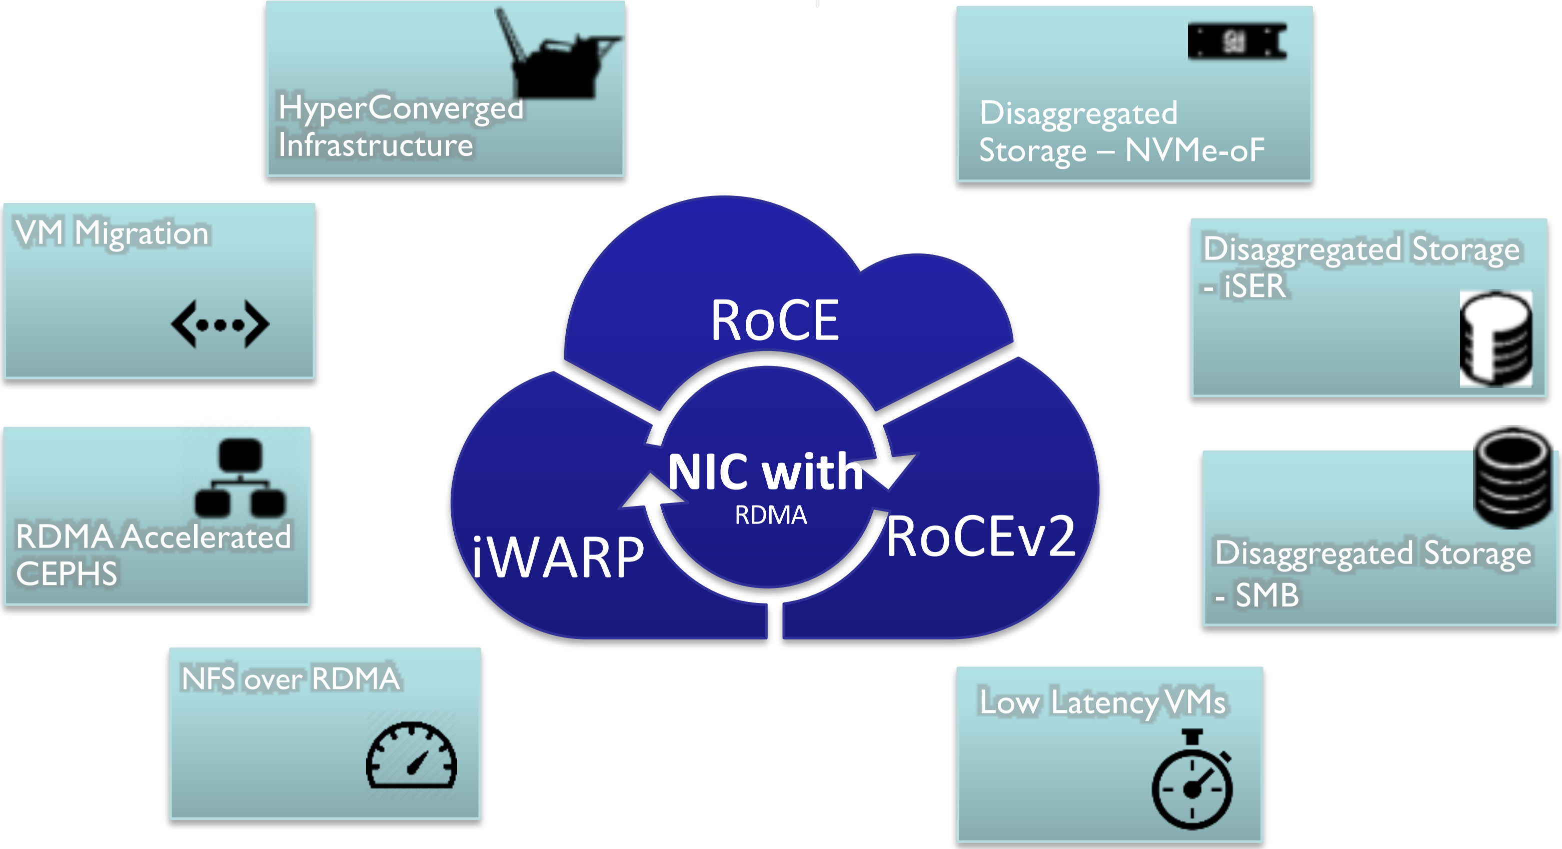

NVMe-oF™ RDMA – potential challenges

Infrastructure and Skillset change?

Congestion	Skillset Requirements	Backward Compatibility
Not Automatic **Not Precise** **Not for everyone**	Keeping the network ‘**lossless**’ RDMA/**OEFD** expertise	RNIC **Upgrade Required** **RDMA Camps** Creates **Islands**

Relationship Status: Microsoft and RoCE

Elden Christensen
@EldenCluster
Follow

After endless support calls with customers struggling with the configuration complexity of RoCE, we have updated our RDMA network recommendations: docs.microsoft.com/en-us/windows-…

Storage Spaces Direct hardware requirements

04/11/2018 • 3 minutes to read • Contributors all

Networking

Recommended (for high performance, at scale, or deployments of 4+ nodes)

- NICs that are remote-direct memory access (RDMA) capable, iWARP (recommended) or RoCE
- Two or more NICs for redundancy and performance
- 25 Gbps network interface or higher

See the Microsoft Blog – comparing the RDMA types
https://blogs.technet.microsoft.com/filecab/2017/09/21/storage-spaces-direct-with-cavium-fastling-41000/

NVMe Transport Performance Comparisons

iSCSI adds 82% more latency, Delivers fewer IOPS

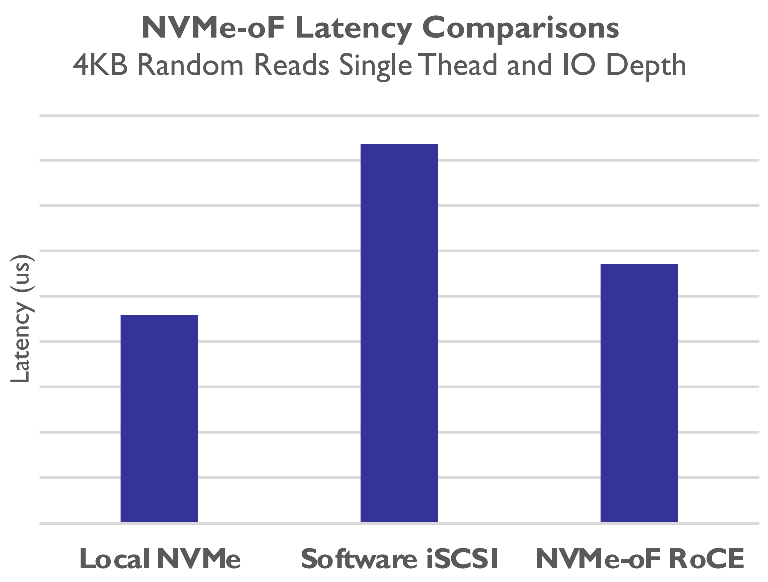

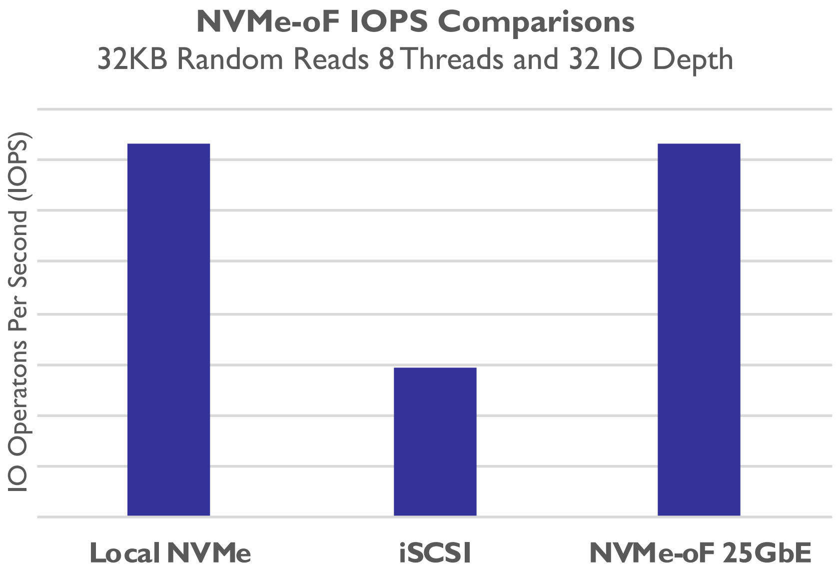

FC-NVMe!

“NVMe” Over Fibre Channel

Transport NVMe Natively over Fibre Channel

Low Latency

Reliable, Secure, Available

Leverage Existing Investments in Fibre Channel

FC-NVMe T11 Committee

Ecosystem Ready

FCP vs. FC-NVMe

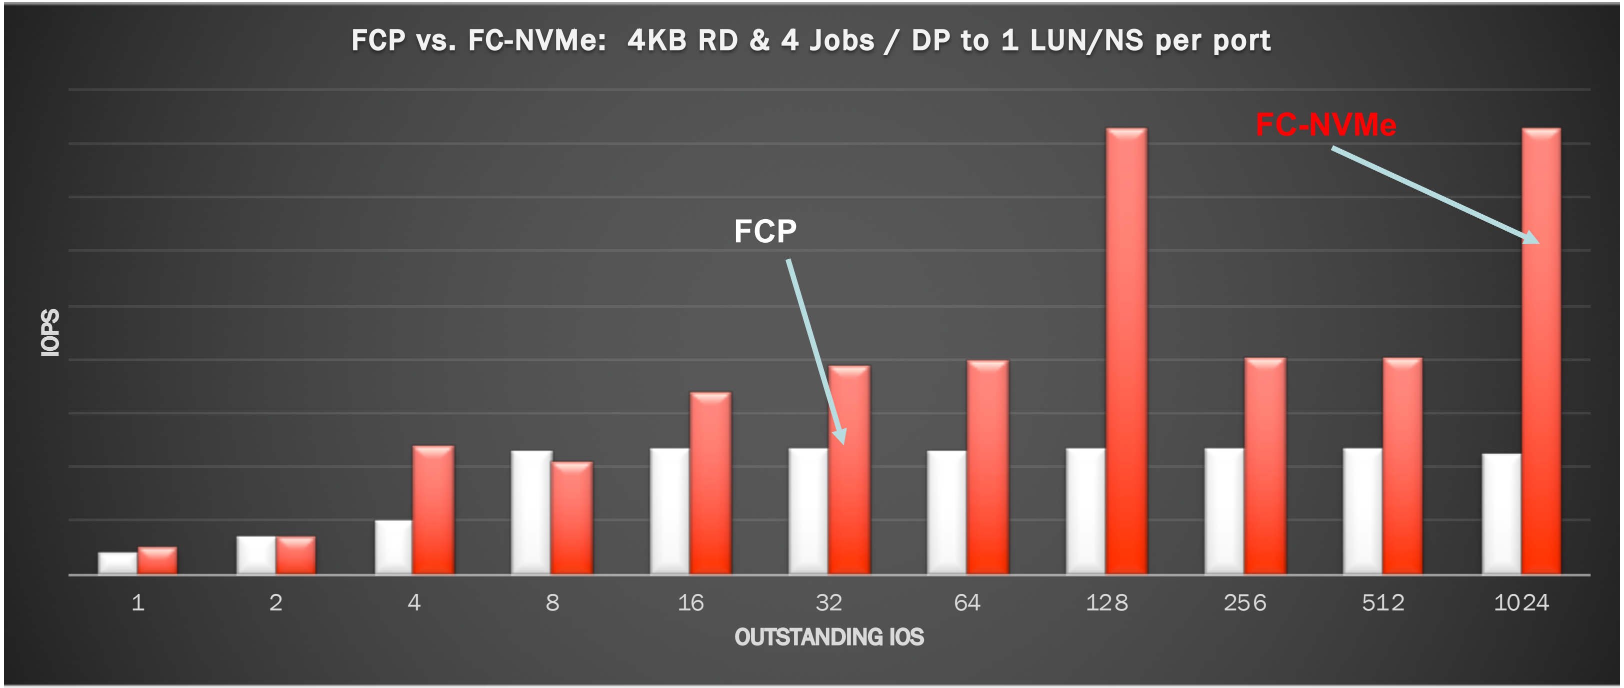

FC-NVMe Scales in performance

Use Cases by Fabric

No one size fits all!

DAS, HPC, AI/ML	Enterprise Applications	All Applications
NVMe/RDMA (Ethernet)	**FC-NVMe (Fibre Channel)**	**NVMe/TCP (Ethernet)**
Performance at the cost of complexity	Leverage existing infrastructure. Reliability is key	Simplicity is key. Balance of performance and cost

Logos are indicative of workload characteristics only.

NVMe/TCP

NVMe-oF: NVMe/TCP

- What: Defines a TCP Transport Binding layer for NVMe-oF
- Promoted by Facebook, Google, Intel, Marvell etc.
- Not RDMA-based, Standardized on 15NOV18
- Why:
  - Enables adoption of NVMe-oF into existing datacenter IP network environments that are not RDMA-enabled

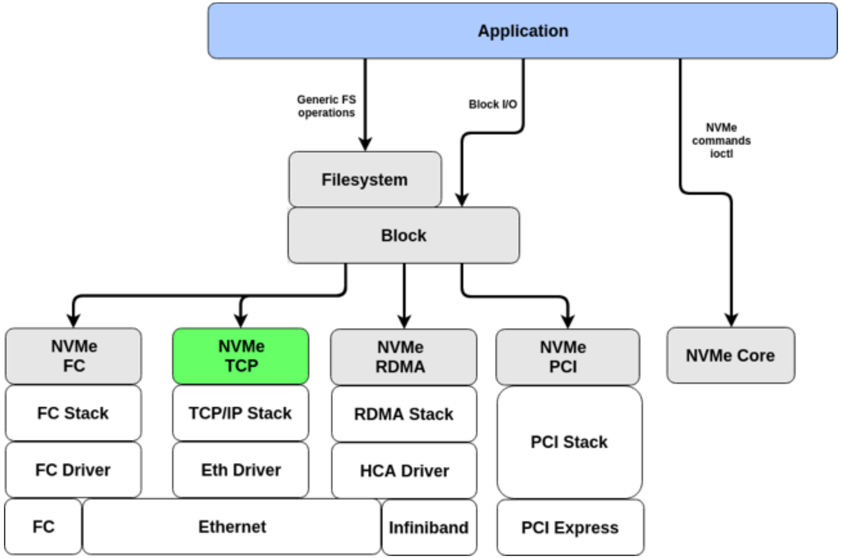

NVMe-oF Driver Stack

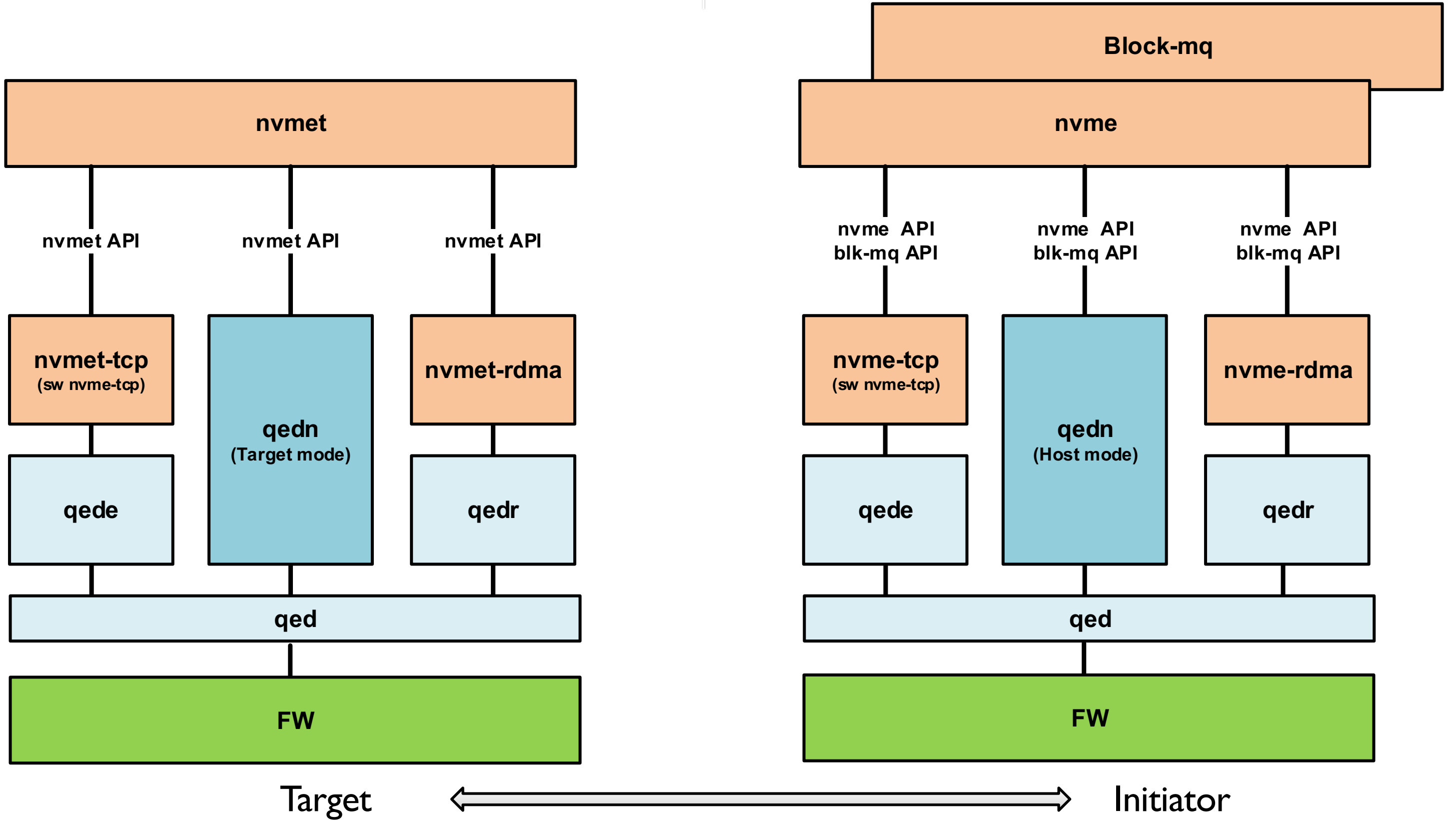

Offloading NVMe/TCP

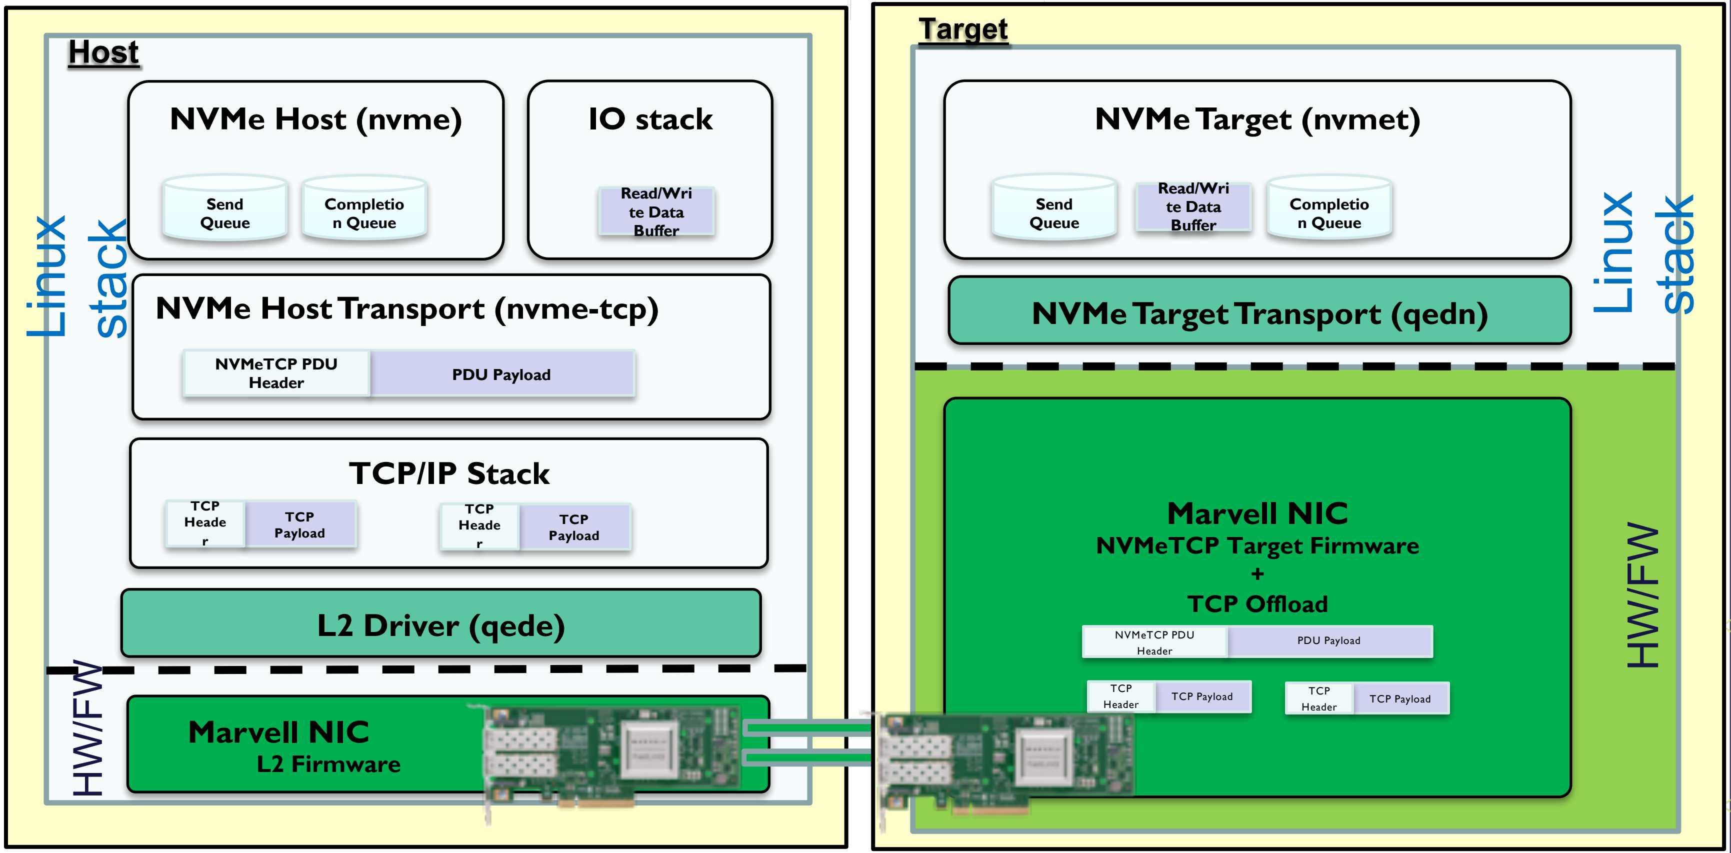

Accelerating NVMe/TCP

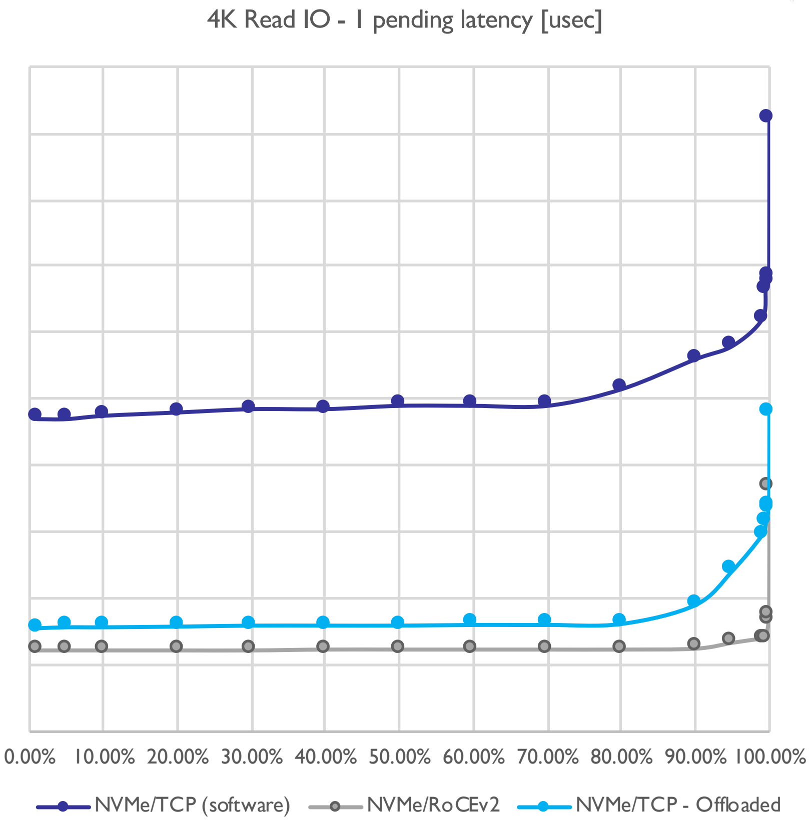

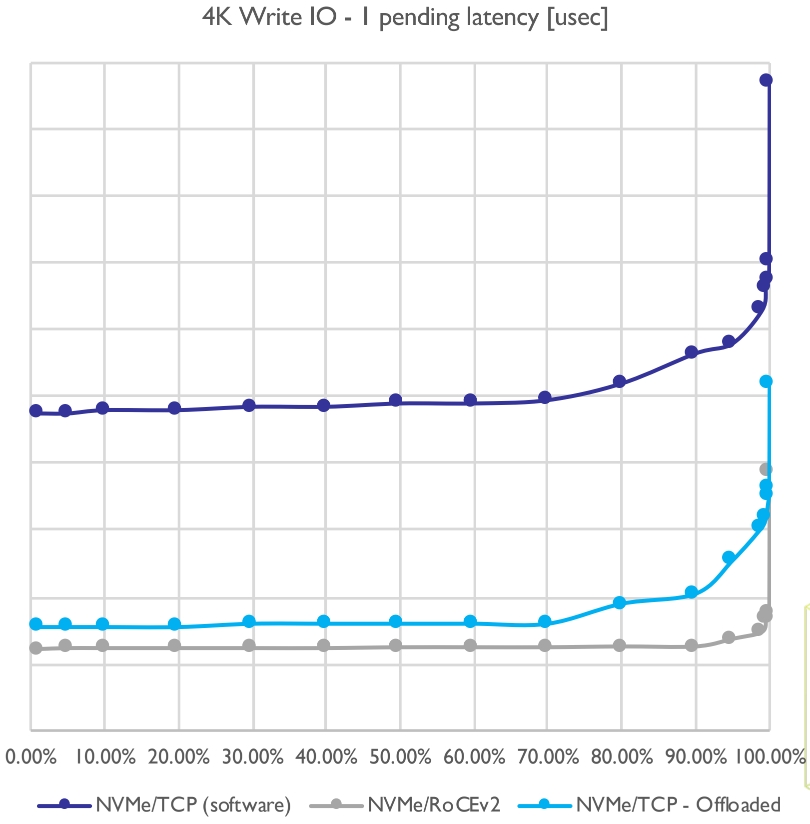

Cost of I/O – NVMe/TCP

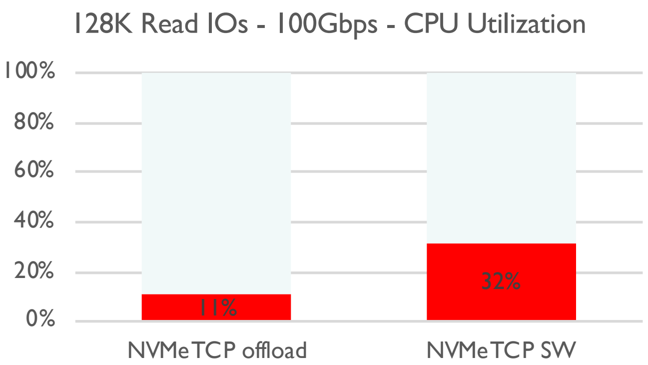

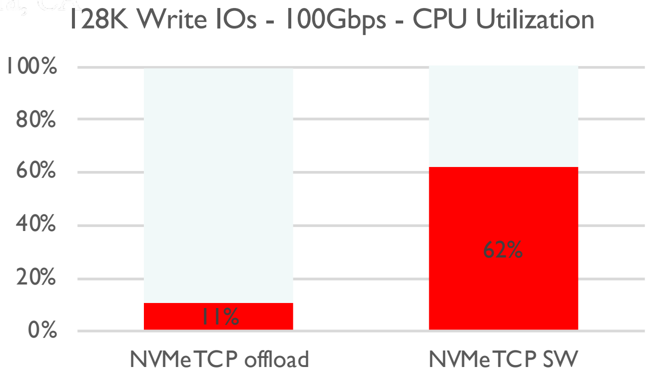

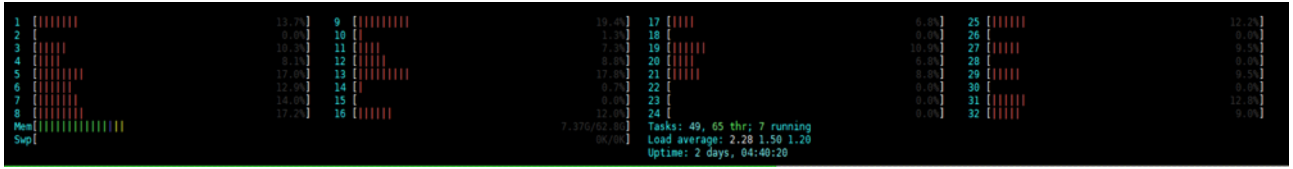

Significant CPU Savings with NVMe/TCP Offload

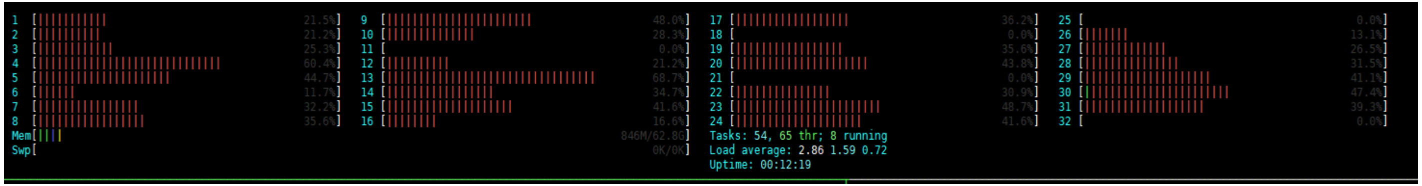

SDC 19

Security FC-NVMe

Drivers for FC-NVMe Security

Security and Privacy Sensitive Verticals

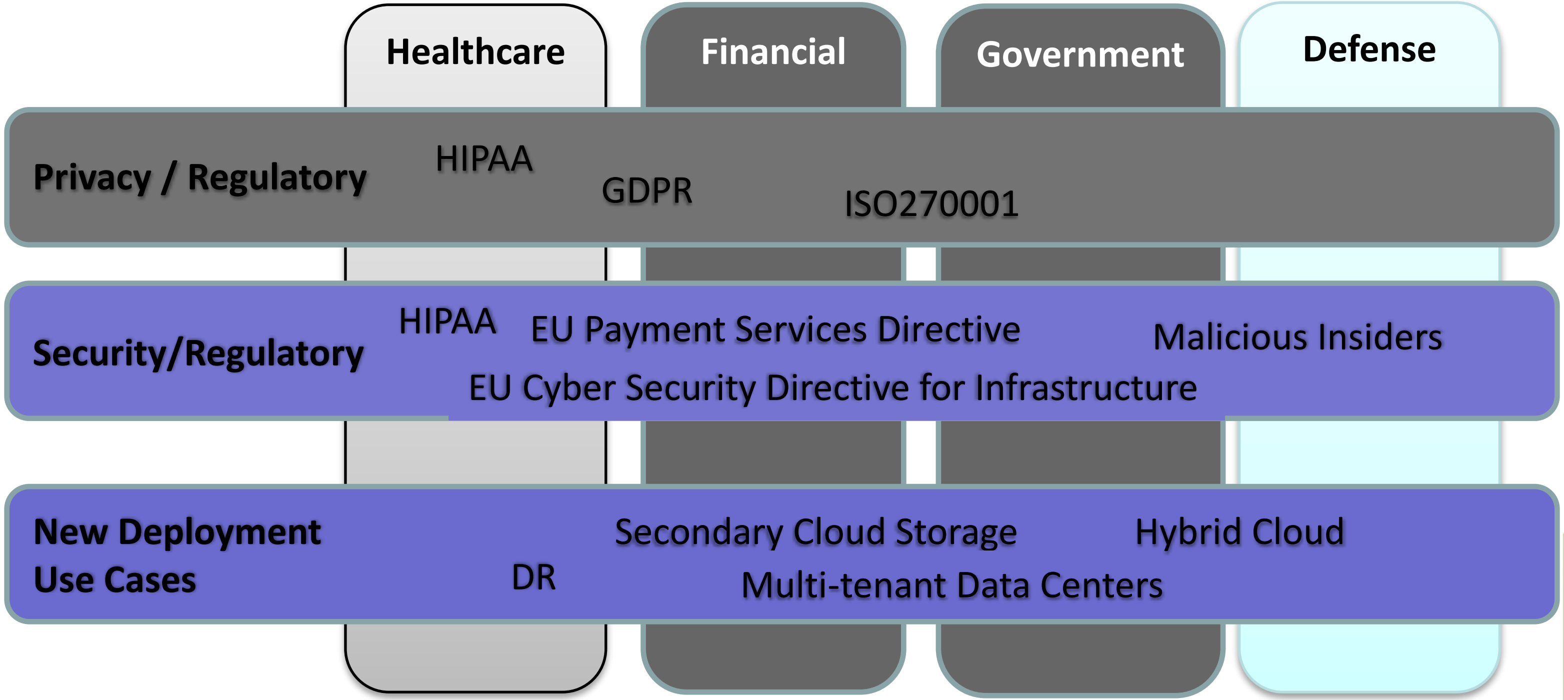

Cost of a data breach and Recent events

Source: IBM Security

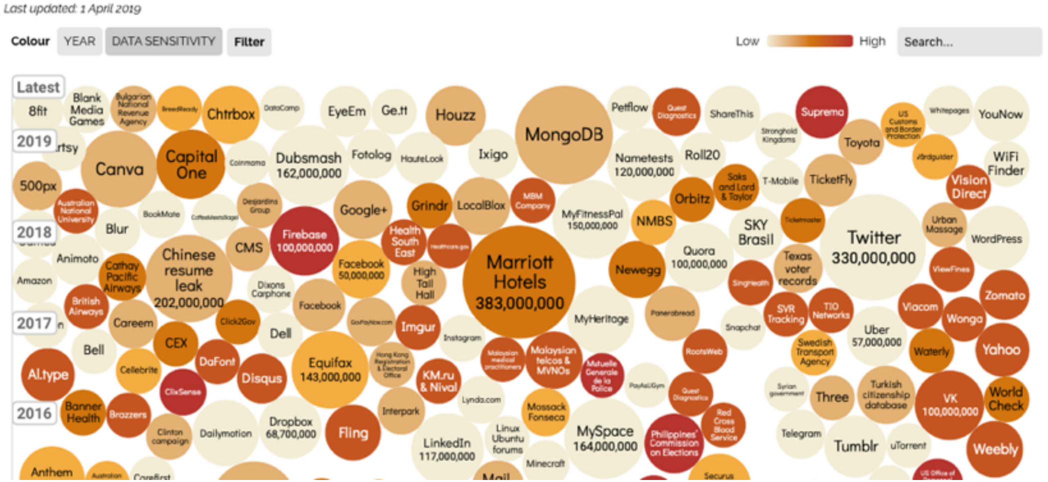

sources: databreaches.net, IDTheftCentre and media reports

None of these breaches have been directly attributed to Fibre Channel

Isn't FC Secure Already?

Trusted Storage Interconnect for Decades

Physical Security	• Data Centers are physically secured
Segregation	• Fibre Channel SANs are segregated networks
Partitioning	• FC Zoning ensures fabric partitioning
Masking	• LUN masking restricts access to specific LUNs
Management	• Out-of-Band Management (IP) is secure, OS Controls

Yes, But…

- New Data Center Architectures bring new threats
  - Distributed data centers - Remote replication and DR backups may be accessed by different users over Fabrics that span several sites
  - Multi Tenant data centers – Need to segregate and protect data traversing the same wire
- Increasing scale of FC SANs
  - Networks can be misconfigured
  - Fabric configuration databases are shared, have WKAs
- Existing mechanisms may not be enough
  - Switches are the sole entity that grant/deny access
    - Authorization based
  - "Segmentation" tools being used to implement "Security"
    - Soft zoning, LUN Masking

Potential DC Storage Security Threats

Sniffing Storage Traffic

Storage Masquerading

Data Corruption

Session Hijacking

Mitigated by Fibre Channel SAN Security

FC-SP-2: What and Why?

- **Why?** : Need to transition SANs from **Authorization and segmentation** based FC security to **authentication and encryption** based security!

- **What?** FC-SP-2 is a ANSI/INCITS standard (2012) that defines protocols to –
  - **Authenticate** Fibre Channel entities
  - **Setup** session **encryption keys**
  - Negotiate parameters to ensure per **frame integrity and confidentiality**
  - Define and **distribute security policies** over FC

- Designed to protect against several classes of threats

Fabric Security Architecture

Components of FC-SP-2 Security Architecture

Authentication Infrastructure	Authentication	Security Associations	Crypto Integrity Confidentiality	Authorization
Secret, certificate, password and pre-shared key based architecture	Protocol to assure identify of communicating entities, negotiation of security requirement and protocol	Protocol to establish Shared key between communicating entities, Based on IKEv2 (RFC4595)	Frame by frame encryption, replay protection, origin authentication, ESP_Header or CT_Authentication	Fabric policies that control which entities can connect with each other, management access to the fabric

FC-SP-2 *ESP_header*

- *ESP_header* (optional) is a layer 2 security protocol that provides
  - Origin authentication
  - Integrity
  - Anti-replay protection
  - Confidentially
- Encapsulating Security Payload (ESP) is defined in RFC 4303
- FC-FS-3 defines optional headers for Fibre Channel, FC-SP defines how to use ESP in Fibre Channel
- Similar protections exist for CT_Authentication

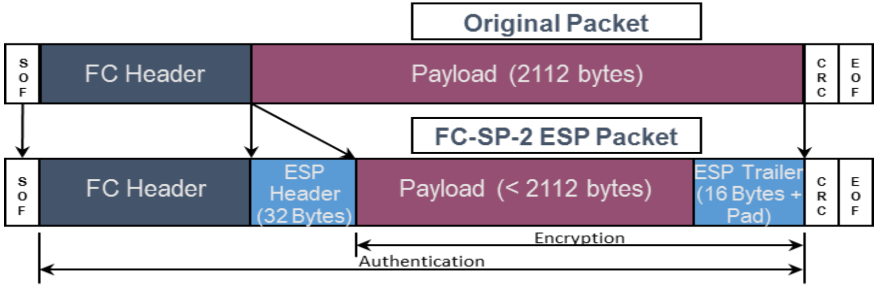

Silicon Root of Trust

Protecting the Integrity of Fibre Channel Firmware

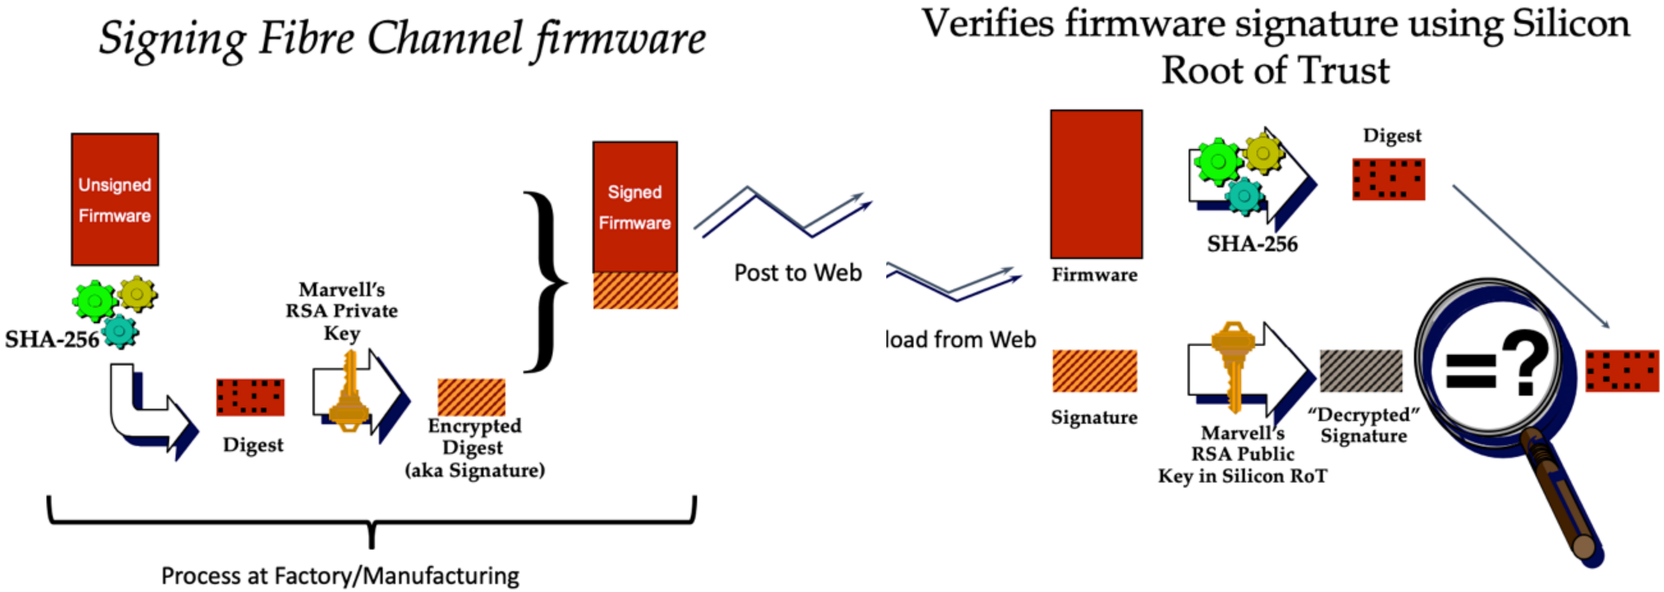

Key Takeaway

Making the right "fabric" choice!

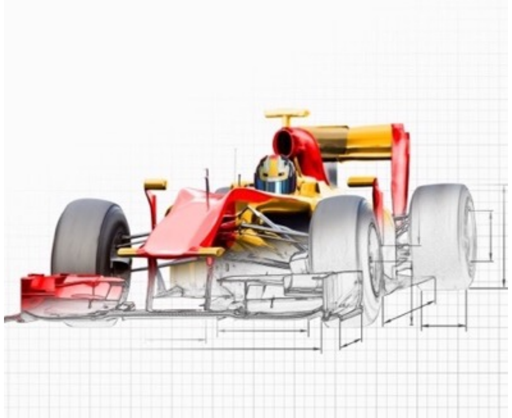

Not "just" about "fabrics" performance

Culture and Install Base

Use Cases and Security

That's it!