Adaptable Storage Acceleration Platforms for Exabyte-scale Data Centers

Manish Muthal
Vice President Datacenter Marketing
August 9, 2018

XILINX

Computing After Moore's Law

- Frequencies have hit a brick wall
- Cost per unit of yielded area is going up
- Compute is becoming a precious resource that must be conserved

Moore's Law Slows While Costs Continue to Increase*

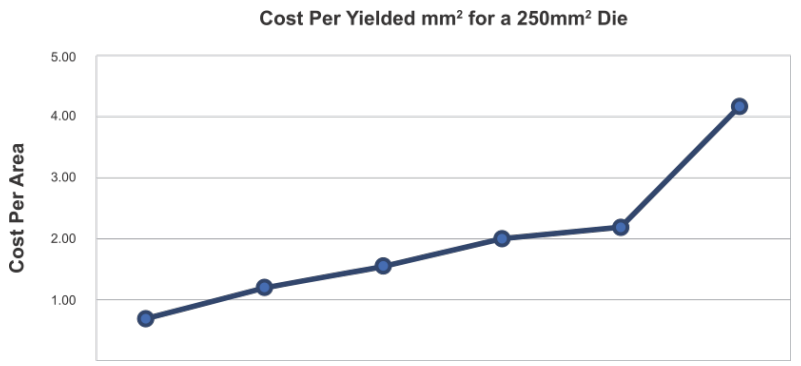

*Source: AMD

Moore's Law is Slowing*

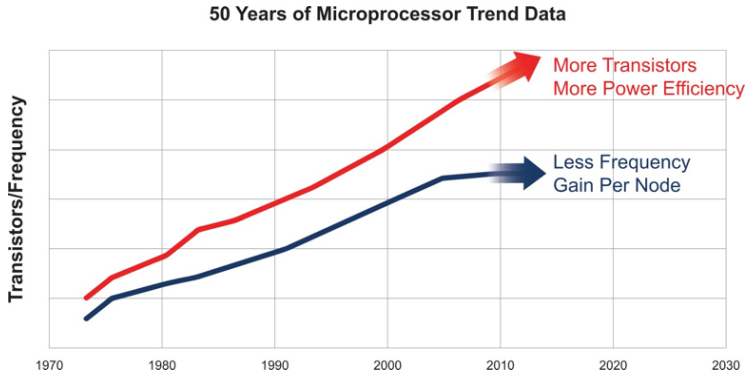

*Source: AMD

XILINX.

Explosion of Data

- Video and IoT driving explosion of data
- Flash Storage driving lower latency and higher bandwidth
- Moving Data Back and Forth between CPU and Storage creates bottlenecks and Consumes Power

More Data from the Network

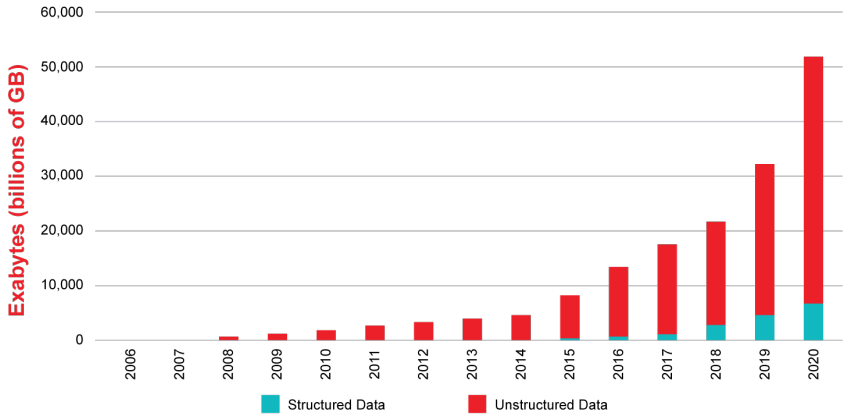

Increasing SSD Densities

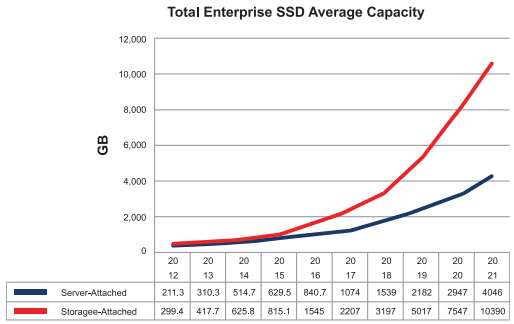

	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Server-Attached	211.3	310.3	514.7	629.5	840.7	1074	1539	2182	2947	4046
Storagee-Attached	299.4	417.7	625.8	815.1	1545	2207	3197	5017	7547	10390

More Data per Server

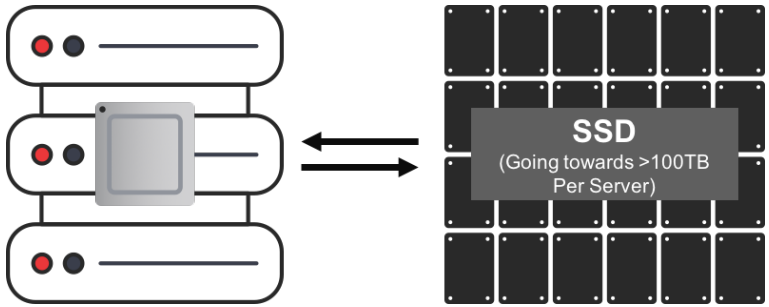

XILINX.

The Dawn of AI

- Machine Learning, Video and Image Processing, Big Data Analytics…
- Increasingly being used in combination
- New Applications and Algorithms coming at a furious pace

Evolution of Machine Learning

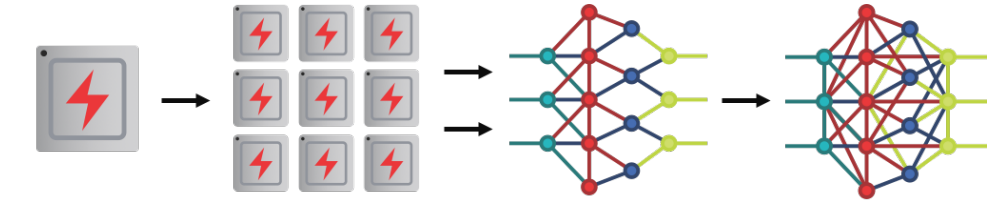

*Source: CellStrat

Explosion of Big Data Apps

Combining Video and Machine Learning

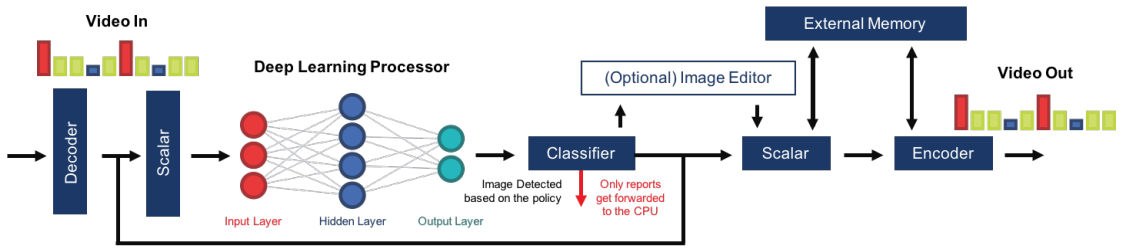

XILINX.

The Answer:
Put Adaptable Acceleration Closer to Storage

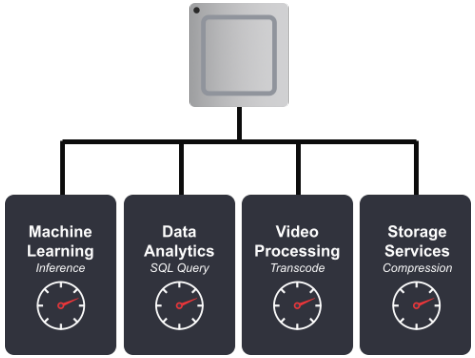

Move Intensive Workloads into Efficient Accelerators

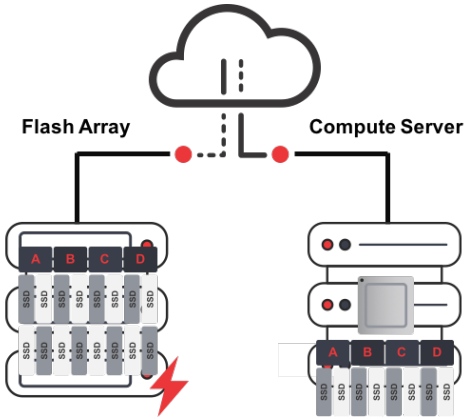

Put Acceleration Closer to Storage Endpoints

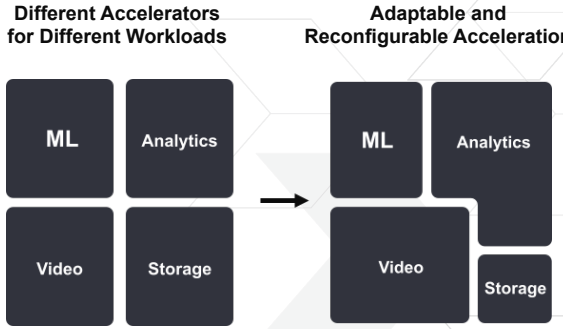

Adaptable Accelerators for Variable and Changing Workloads

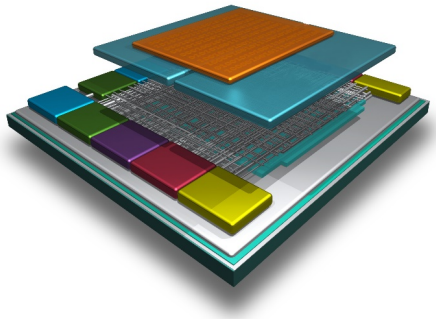

FPGAs Provide the Ideal Platform for Distributed Storage Acceleration

11-40X Improvement in Efficiency (Perf/Watt) over CPU

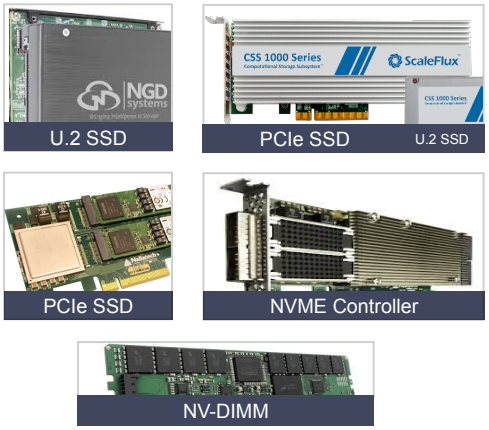

Easy to Attach to a Variety of Storage Endpoints

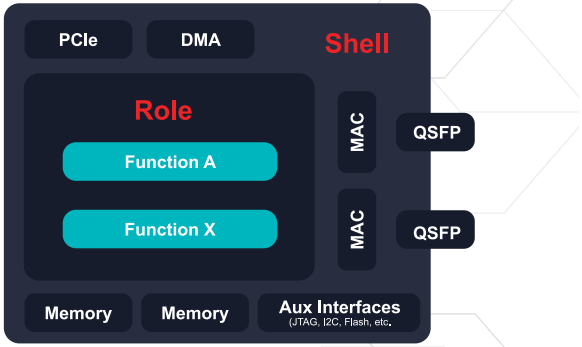

Adaptable and Dynamically Reconfigurable

XILINX.

Solving The Needle In the Haystack Problem

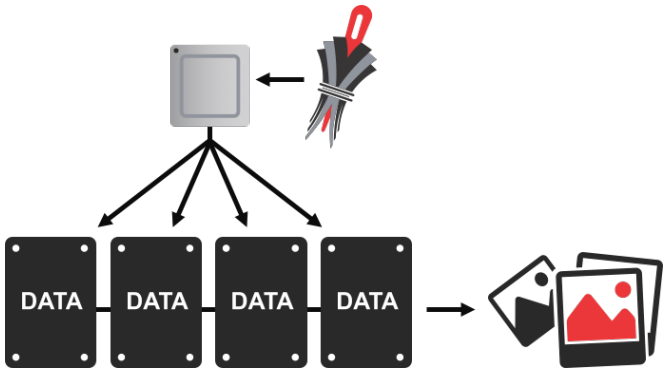

- Search for Needle in CPU/Accelerator
- Sequential Scan of All Drive Data
- I/O and Processing Bottlenecks
- High Power Consumption to Move All Data

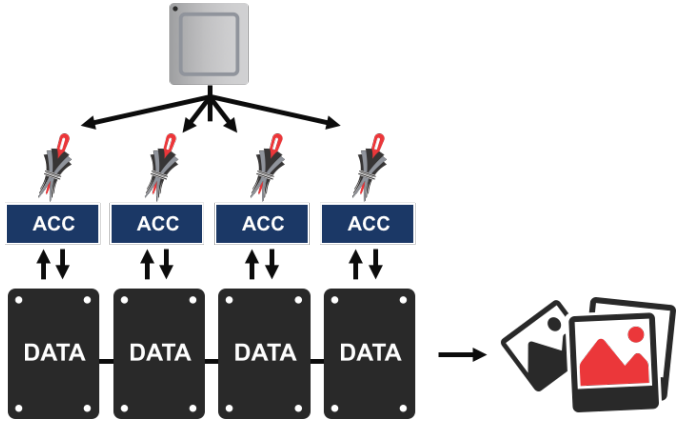

- Move the Needle down to Distributed Accelerators
- Individual Scan from Distributed Accelerators
- Eliminate I/O and Processing Bottlenecks
- Localized Data Inspection Reduces Power

Improving Total Cost of Ownership

Significant CAPEX/OPEX Reduction

- Accelerators free up general compute cycles
- Result is more work per server or less servers per work
- Accelerators add only incremental cost/power (10-20%)
- Results in Capex/Opex gains for DC Operators

Server Cost = X
Server Cost Accelerated = X + 15%
Server Power = Y
Server Power Accelerated = Y + 20%

		Server CPU Utilization	
DC Workload Mix	**Acceleration Ratio**	Unaccelerated	Accelerated
General Purpose Compute	N/A	40%	78%
Storage Services	40	10%	0.25%
Database/Big Data	34	10%	0.29%
Video Transcoding	40	10%	0.25%
Machine Learning	11	10%	0.91%
Idle/Overhead	N/A	20%	20%

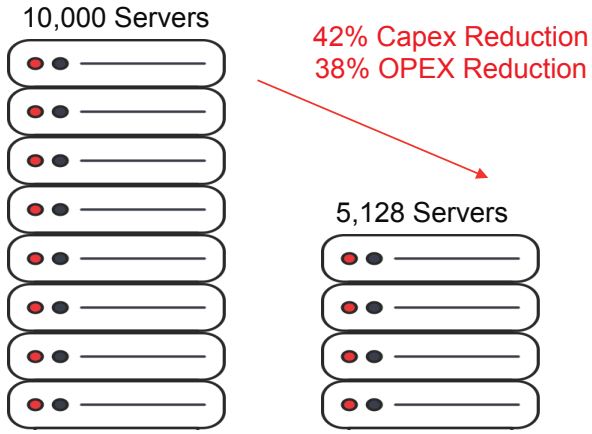

XILINX.

The Application Ecosystem for the Adaptable Storage Platform is Strong and Growing Rapidly

DEVELOPERS

100%
QoQ Growth of Published Applications in FY18

Hundreds
of Developers Trained Every Quarter

1500+
Developers Trained by Year End

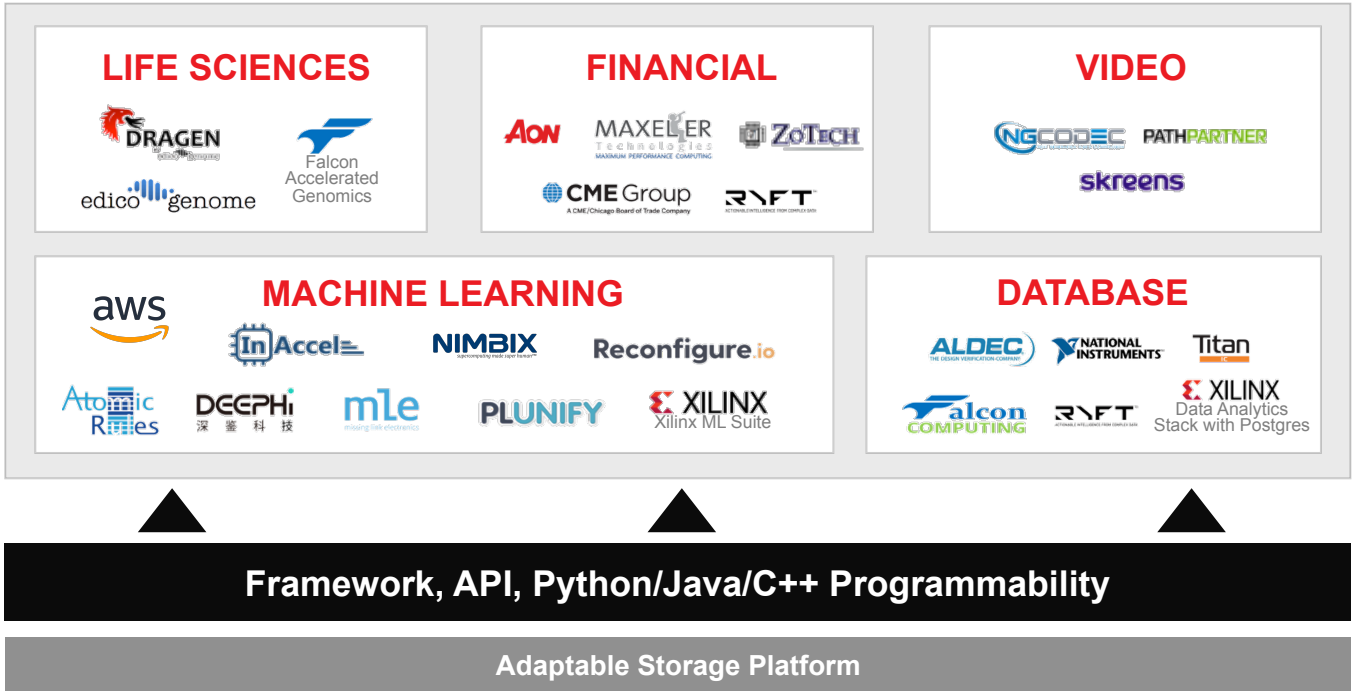

XILINX.

7nm Rocket Fuel

- New Device Category for Adaptive Workload Acceleration
- Network-On-Chip
- HW/SW Programmable Engines
- Coming to Adaptable Storage in 2019!

Xilinx ACAP Architecture

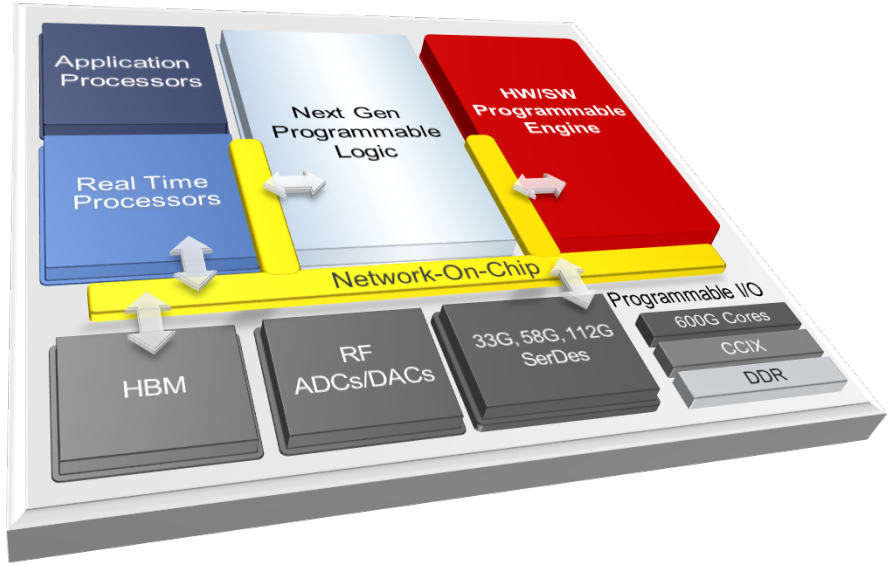

XILINX.

Thank You

- Please Come Visit Xilinx in our Booth #313
- Live Demos!
  - Storage Compression Offload
  - Programmable Controller for Software-defined Flash
  - Database Acceleration on FPGAs
  - NVDIMM-N Solution

Adaptable.
Intelligent.

XILINX