

An abstract graphic on the left side of the image, composed of numerous thin, light green lines that curve and swirl together to form a complex, organic shape resembling a stylized flower or a tunnel. The lines are set against a solid dark blue background.

Open. Together.



OCP
SUMMIT

Bringing Intelligence to Storage with SmartSSD

Kristian Vättö, SSD Product Marketing
Samsung Semiconductor Europe GmbH



Moving Data is Easier than Before or not

1956

5MB on a trailer



2016

100PB in a truck



Future



10EB+ where???



Bandwidth is Not Catching Up

[Normalized]

1000x

100x

10x

1x

'00 '01 '02 '03 '04 '05 '06 '07 '08 '09 '10 '11 '12 '13 '14 '15 '16 '17 '18 '19 '20

Starting points

Interface: Ultra ATA 133

Capacity: 80GB

— Interface [MB/s] — Capacity [TB]

SATA 3.0

PCIe Gen3x4

PM1733 PCIe Gen4x4

PM1633 16TB

PM1643 32TB

376x increase in 20 years

Capacity grew 8x faster

48x increase in 20 years

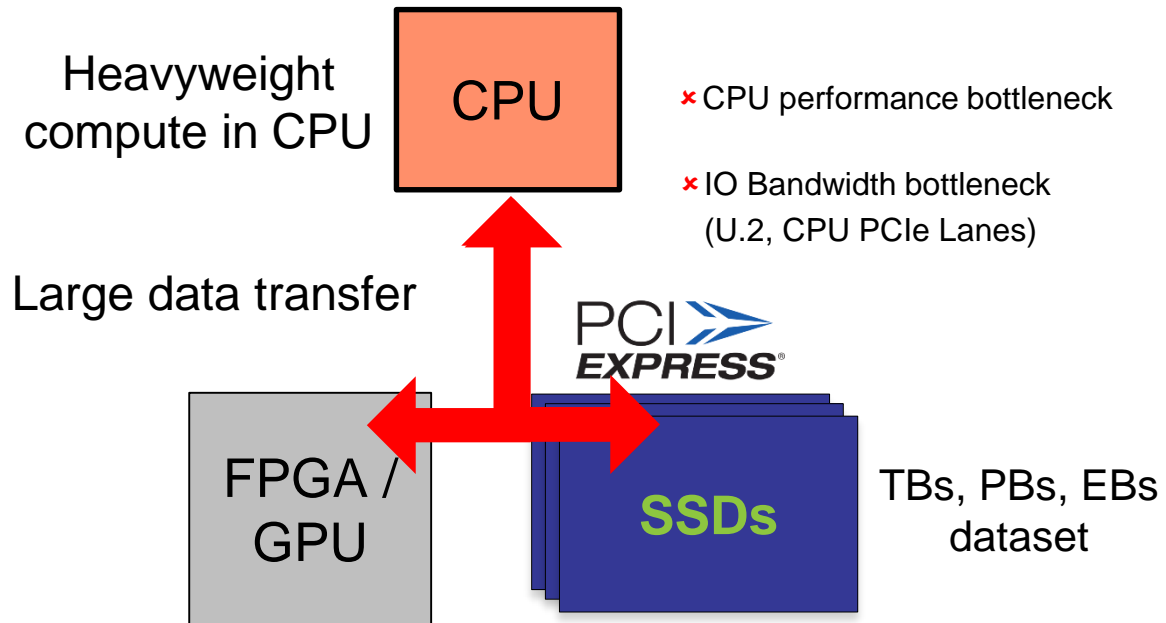
Sources

Interface: Company press releases

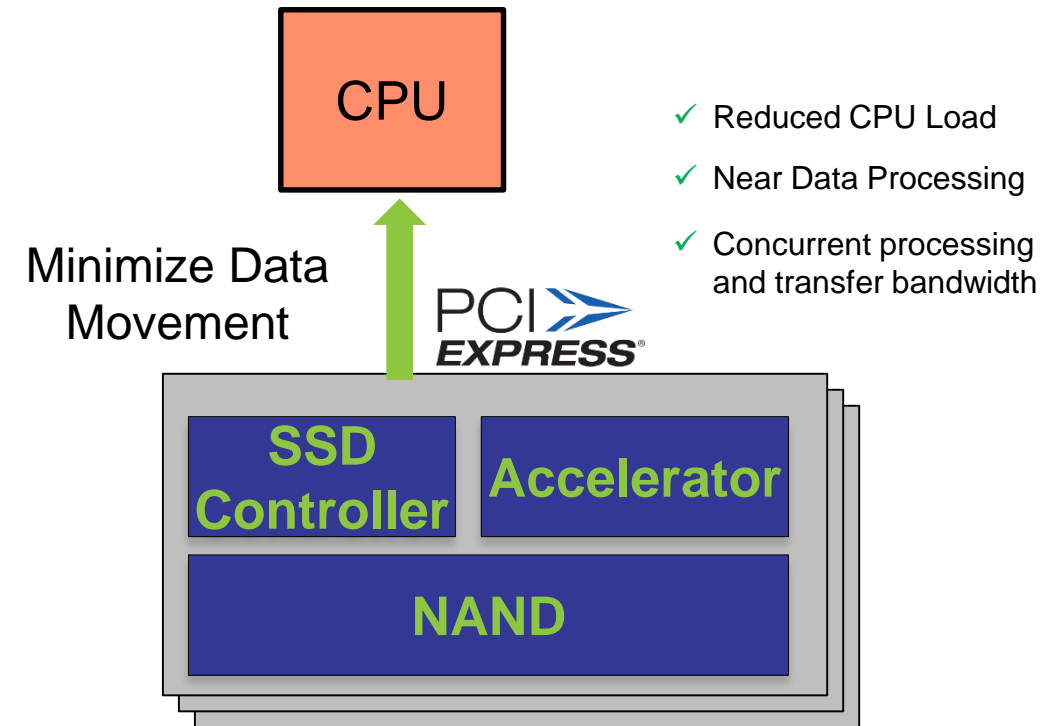
Capacity: Gartner & press releases

SmartSSD: Computational Storage

Today's Architectures




Storage and Data Acceleration



Proving the SmartSSD Concept


PM983F AIC

PoC using PCIe add-in card
Successfully integrated with Spark/Bigstream
Easily ported several data-intensive workloads

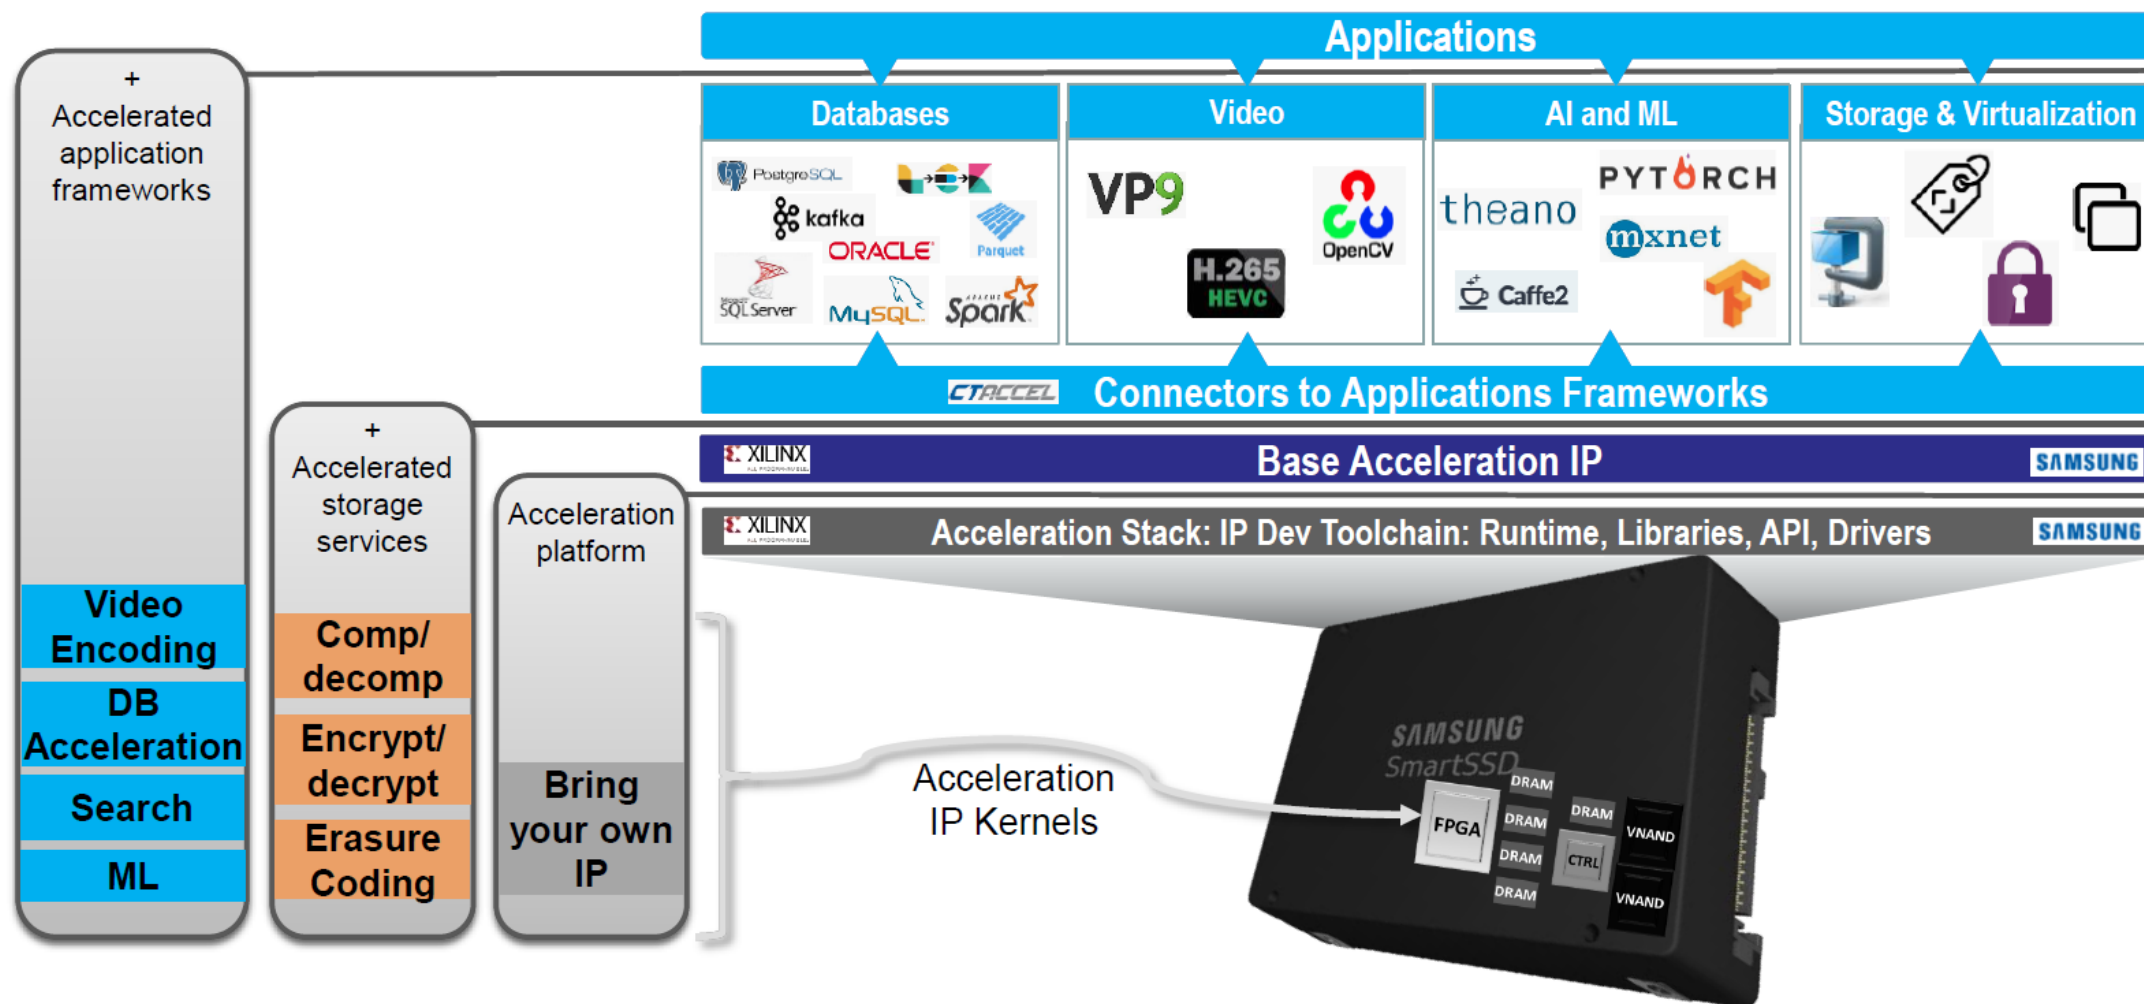


PM983F Results

For I/O-bound workloads, SmartSSD showed 3x to 4x better performance, with scalability



SmartSSD Platform & Ecosystem



I/O bound tasks

- 2 choices:

only to discard most of it, reduce it down to a scalar, or worse still write it all back only slightly modified

Diagram illustrating the Reference Architecture for a data warehouse, showing the flow of data and processing components.

Data @rest (Input/Output) is connected to the processing stack.

Processing Stack (Left):

- parse
- index
- stats
- compress
- encrypt

Processing Stack (Right):

- Scan-filter
- parse
- decompress
- decrypt

Callout: Pack an entire DB storage engine into each drive

SERVER

Reference Architecture

Save expensive server CPUs for monetizable DWUs, DTUS

Unlimited Concurrency

Dense Storage Nodes

Density not frontend-limited

- Not just compression and encryption
- Offload entire threads

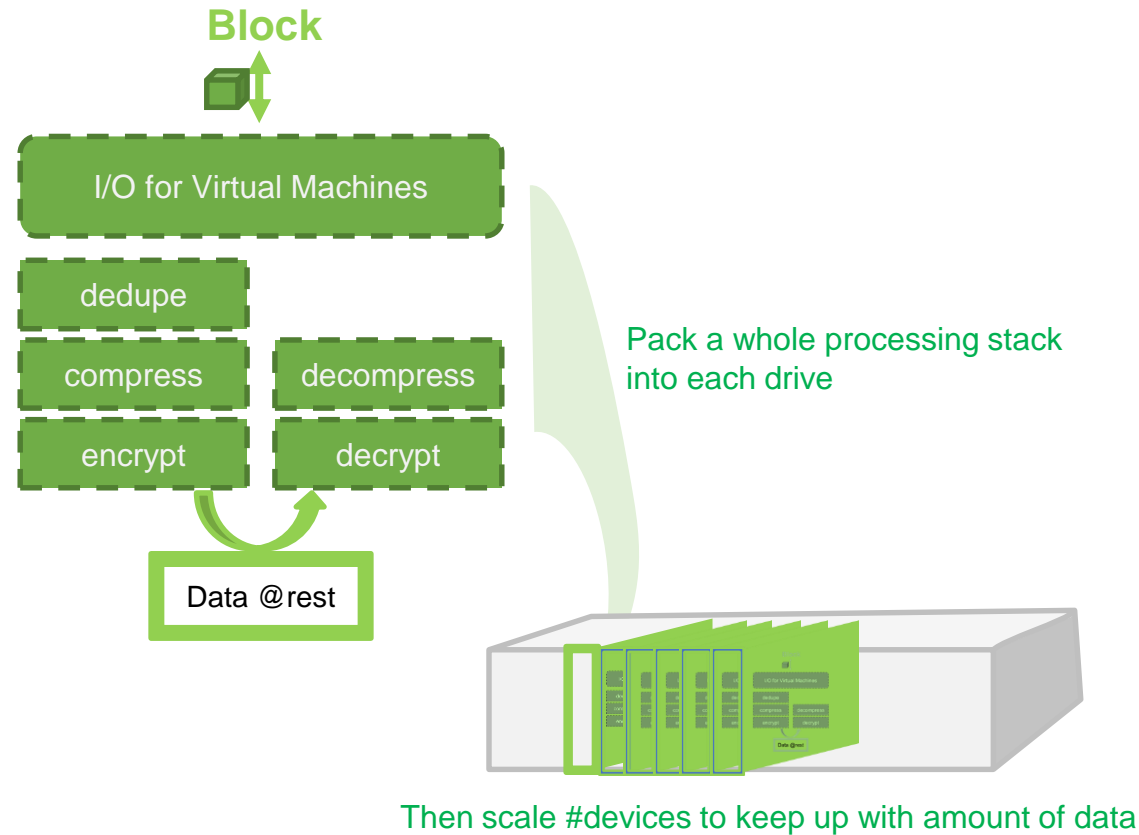
Scaling to high device counts

A. Processing capacity

Enough processing and memory in each device to handle its own work of compaction and dedup


B. Data Bandwidth

Each device has internal data path to processing
Net data processing bandwidth scales linearly



SmartSSD Technology Roadmap

Roadmap to smaller FF (U.2) and greater integration with SSD controller




SmartSSD Acceleration Use Cases



Financial Services
Fraud Detection



Video Transcoding



Immediate insights
from Data lakes



Instantaneous and multi-rules
Fuzzy search



Real-time Log Analytics



Speedup Complex
Ad-hoc queries



The only limit is your imagination


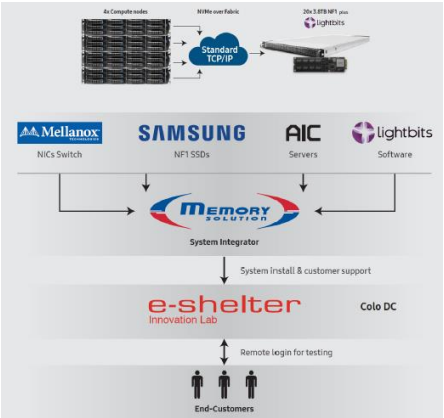
Come see us at booth B7!



PM1733/35
First PCIe Gen4 SSD
Now available!

NVMe over TCP
Remote PoC testbed

with  lightbits



SmartSSD
Computational Storage

Samsung V-NAND
Samsung Controller
Xilinx FPGA

Key-Value SSD for DB acceleration

