



A New Generation SSD for Big Data Applications

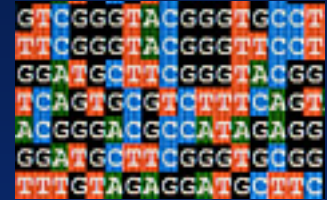
Anthony (Tony) Lavia

Board Director

Cognitive Silicon Group


Big Data needs Smart Storage

- Why now:
 - Massive quantities
 - Geographically distributed
 - Real time
 - Unstructured
 - Visual content



```
GTCGGGTACGGGTGCCT  
TTCGGGTACGGGTTCC  
GGATGCTTCGGGTACGG  
TCAGTGCGTCTTTCAGT  
ACGGGACGCCATAGAGG  
GGATGCTTCGGGTGCGG  
TTTGTAGAGGATGCTTC
```

Big Data needs Smart Storage



- Imagine:
 - Store file IF not already there
 - Save post IF reflects client negativity
 - Retrieve image IF it shows a defective part
 - Retrieve audio file IF contains glass breakage
 - Save video frame IF a person is identified
 - Compress and save video IF suspicious activity

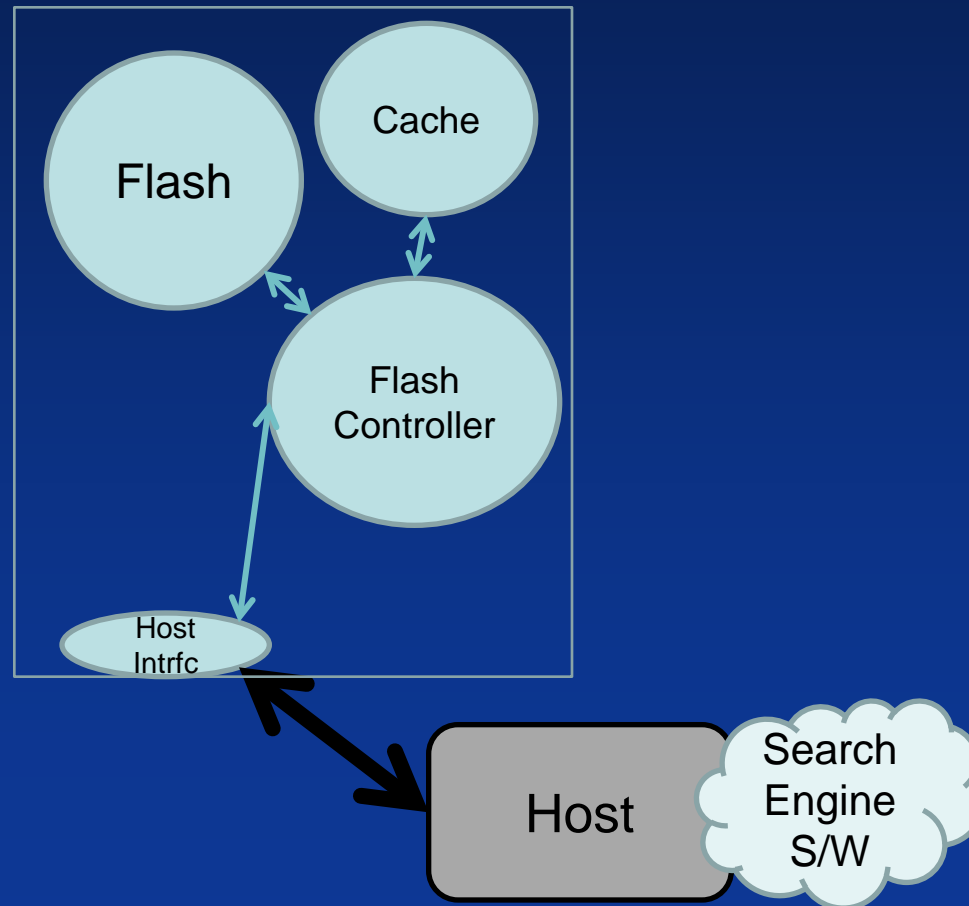


Big Data needs Smart Storage SSD

- Manage Flash for life and reliability:
 - In-line de-duplication
 - Map Blocks/Files
- Amdahl's Law \leftrightarrow Moore's Law
 - Limits of multiple processors

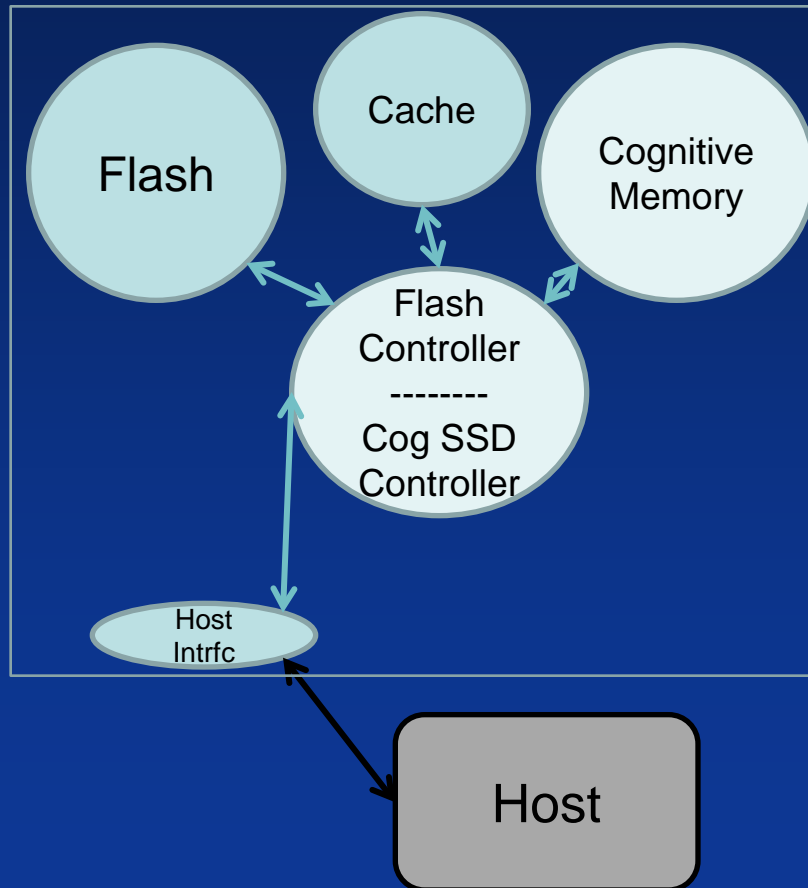
The Concept...

1. Today's SSD




...The Concept

2. Smart SSD



The Enabling Technology Cognitive Memory

- Pattern-matching chip
 - Massively parallel operation
 - In memory processing logic
 - Latency independent of size
 - Exact and probabilistic matching
 - High speed – low latency
 - Low power
 - Moore's Law amenable



Is It Real?

IBM ZISC



CM1K





The Value Proposition...

1. In-Memory Processing

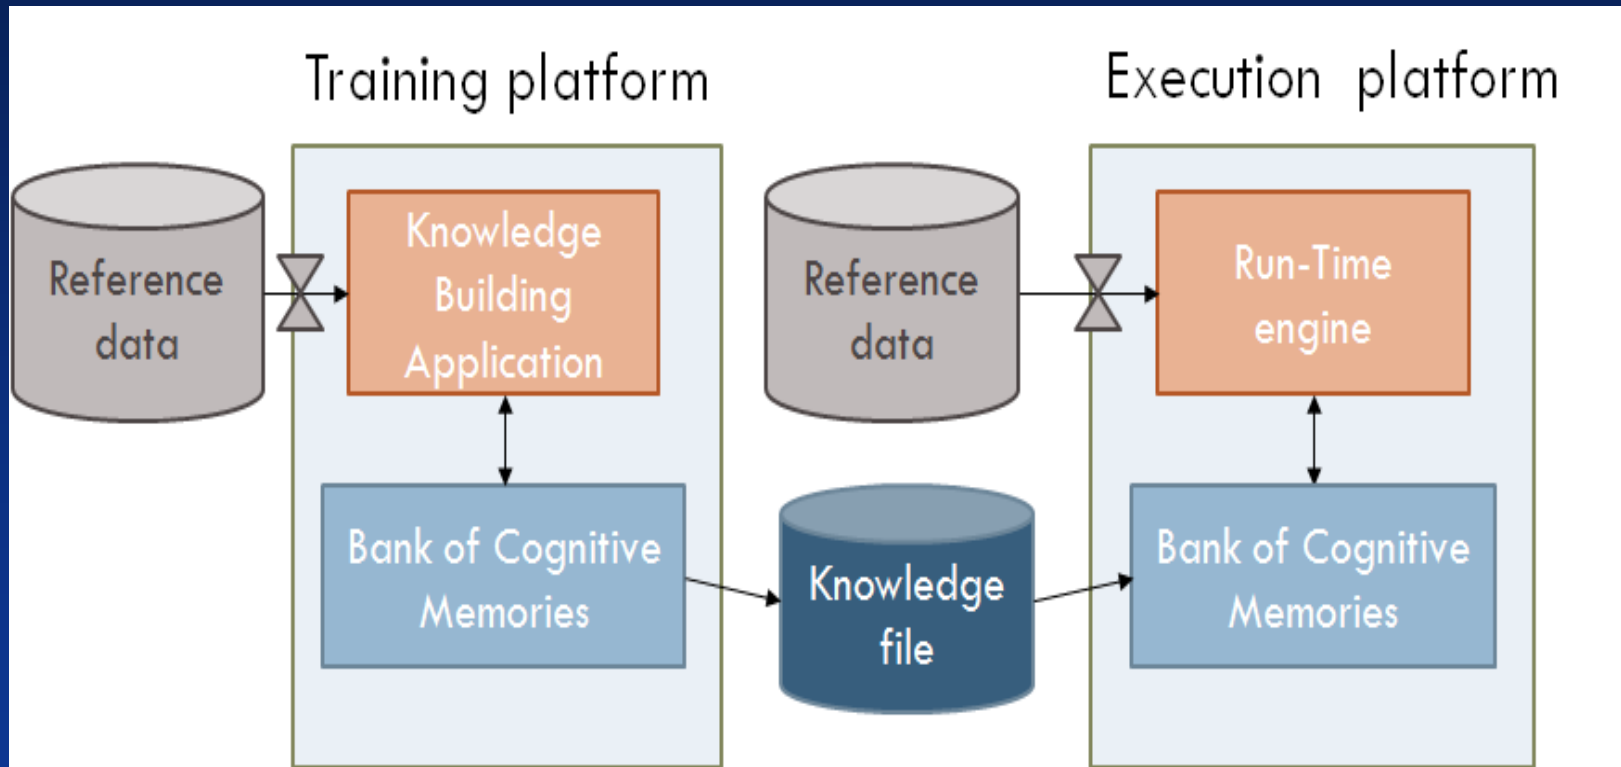
- H/W-accelerated searching
- Off-loads host server
- Scales practically w/o limit
- Super computer performance



...The Value Proposition

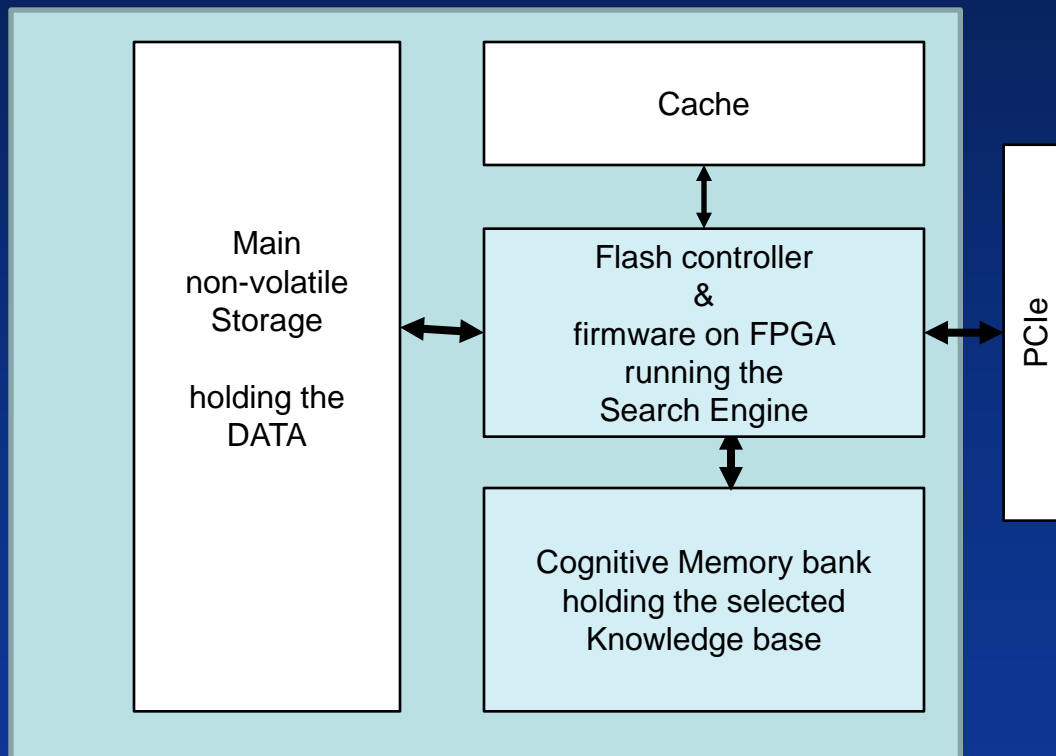
2. Ultra-low power draw
3. Multi Application-specific
 - Various knowledge bases
4. Fuzzy Logic matching
5. Learning capability
6. Train vs Program

Programming: Training vs S/W



Collaboration with Host Server

Smart SSD



Functional Interplay

- Host interface protocols
 - Driver protocol (NVME)
 - DB query overlay cmds
- Cognitive H/W
 - Driver protocol (NVME)
 - DB query overlay cmds
- Controller+
 - Intercept DB cmds
 - Utilizes cognitive memory
 - Search Engine function

The Enabling Technology Today's IC

■ NeuroMem chip



- Introduced in 2011
- 1K cells/chip
- 256B vector/cell
- 10 μ sec response
- <0.5 W/chip
- 27 Mhz clock
- 130 nm

Sample Application Text Analytics

- CRM: Sentiment Analysis
 - Text matching
 - Library of phrases
 - 15M postings to search
 - Benchmark results:
 - 8X chips → 5K transactions/s
 - vs 36 blade servers
 - Consumes 1000X less power
 - Costs much much less



Conclusion

- Big Data = Big Opportunity
- Stretches conventional storage
- With Smart SSD approach:
 - Huge leap forward
 - With incremental technology steps



Thank You Kindly
For Your Attention

Anthony (Tony) Lavia
Board Director
Cognitive Silicon Group