

Ultra Massive application storage capacity for real-time applications

A True Limitless Capacity Data Grid

Shay Hassidim , Deputy CTO ,
GigaSpaces

About myself – Shay Hassidim

- GigaSpaces Deputy CTO
- More than 20 years of SW development
 - Java , .Net , C++ , you name it...
- 11 years with GigaSpaces
 - From the time we have been startup at the garage...
- Oversight the entire technical operations within the Americas
- Located in GigaSpaces NYC HQ



GigaSpaces Business Card

SINGLE PLATFORM. COMPLETE SCALABILITY.

Industry's only application virtualization platform that enables end-to-end scaling with a single product



Financial Services eCommerce



eGaming



Travel & Trans.



Telco



Healthcare

- 300+ Direct Customers
- 75+ Cloud Customers
- 30,000 Sites running GigaSpaces
- 25+ ISVs

Founded 2000

Selected Customers...

| | | | | | |
|---|---|---|---|---|---|
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |

Speed




Scalability



Simplicity

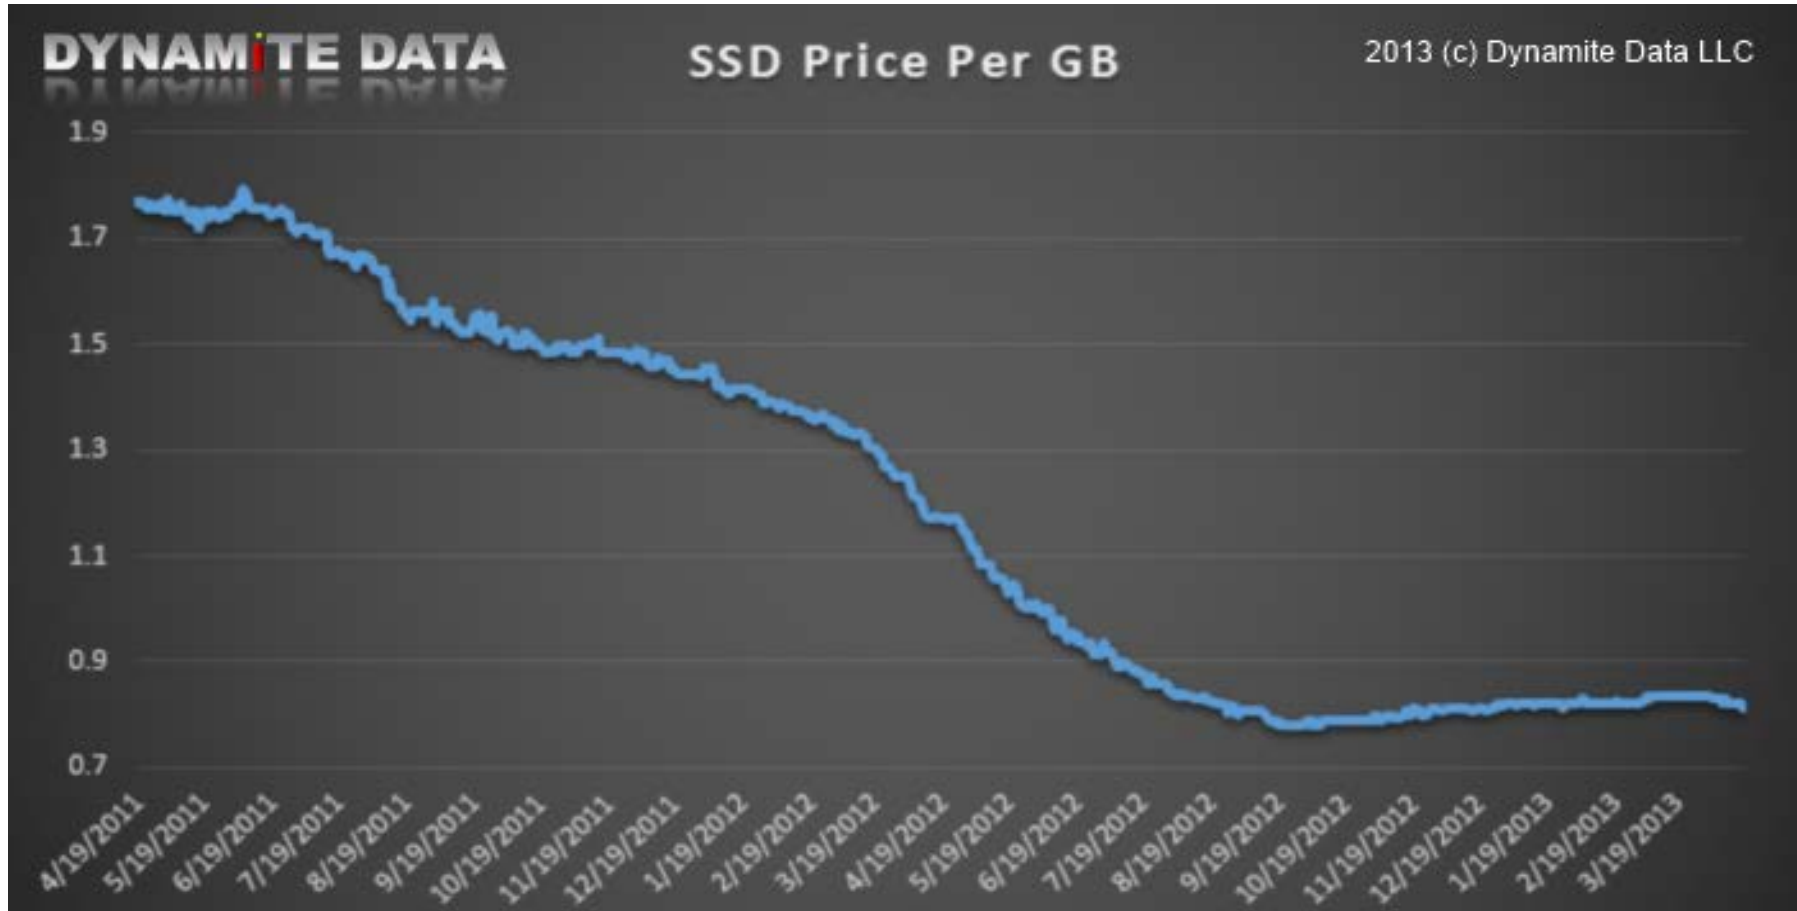


What's your largest in-memory compute data grid?






SSD is Everywhere - Big Time and it's Not That Expensive!



The Dream-Team Combination



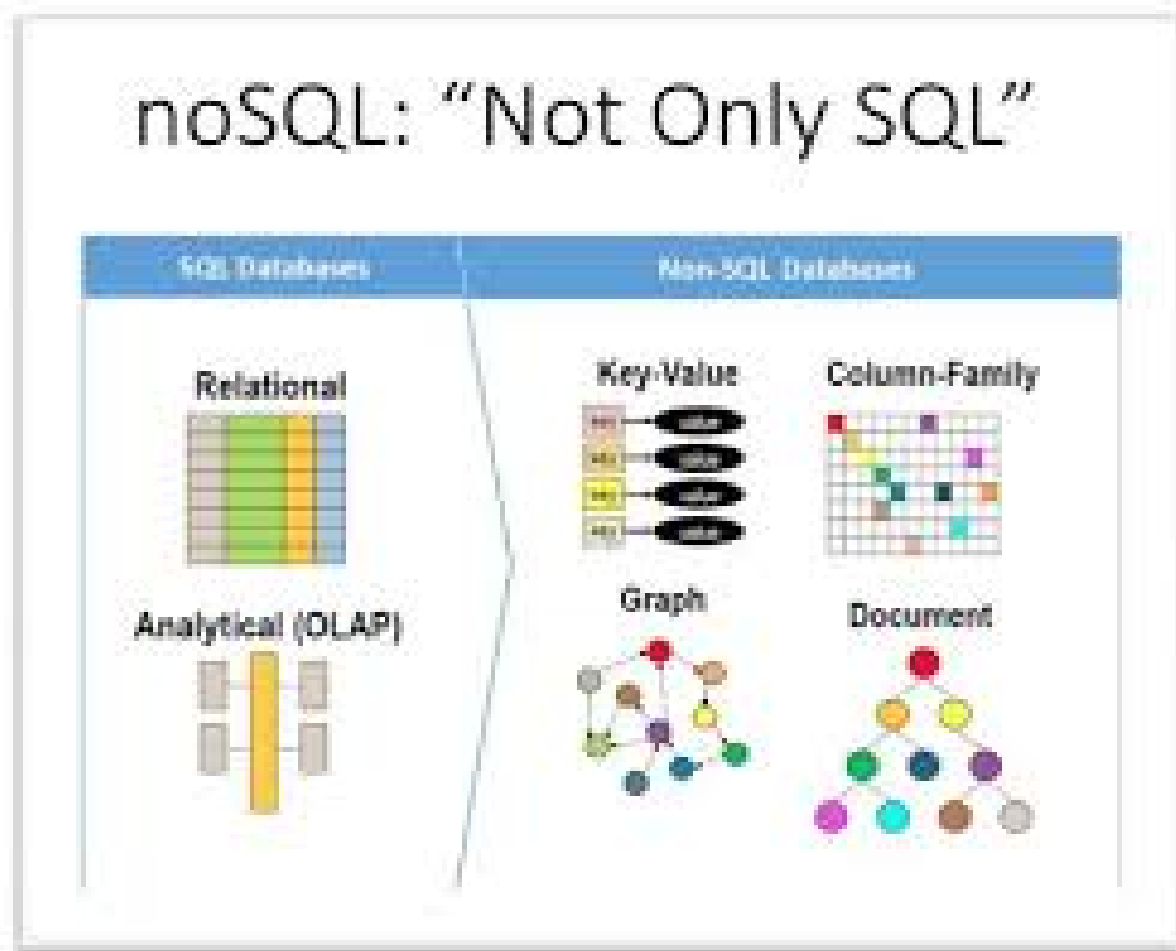
What's the Motivation Moving into a SSD-based Data Grid?



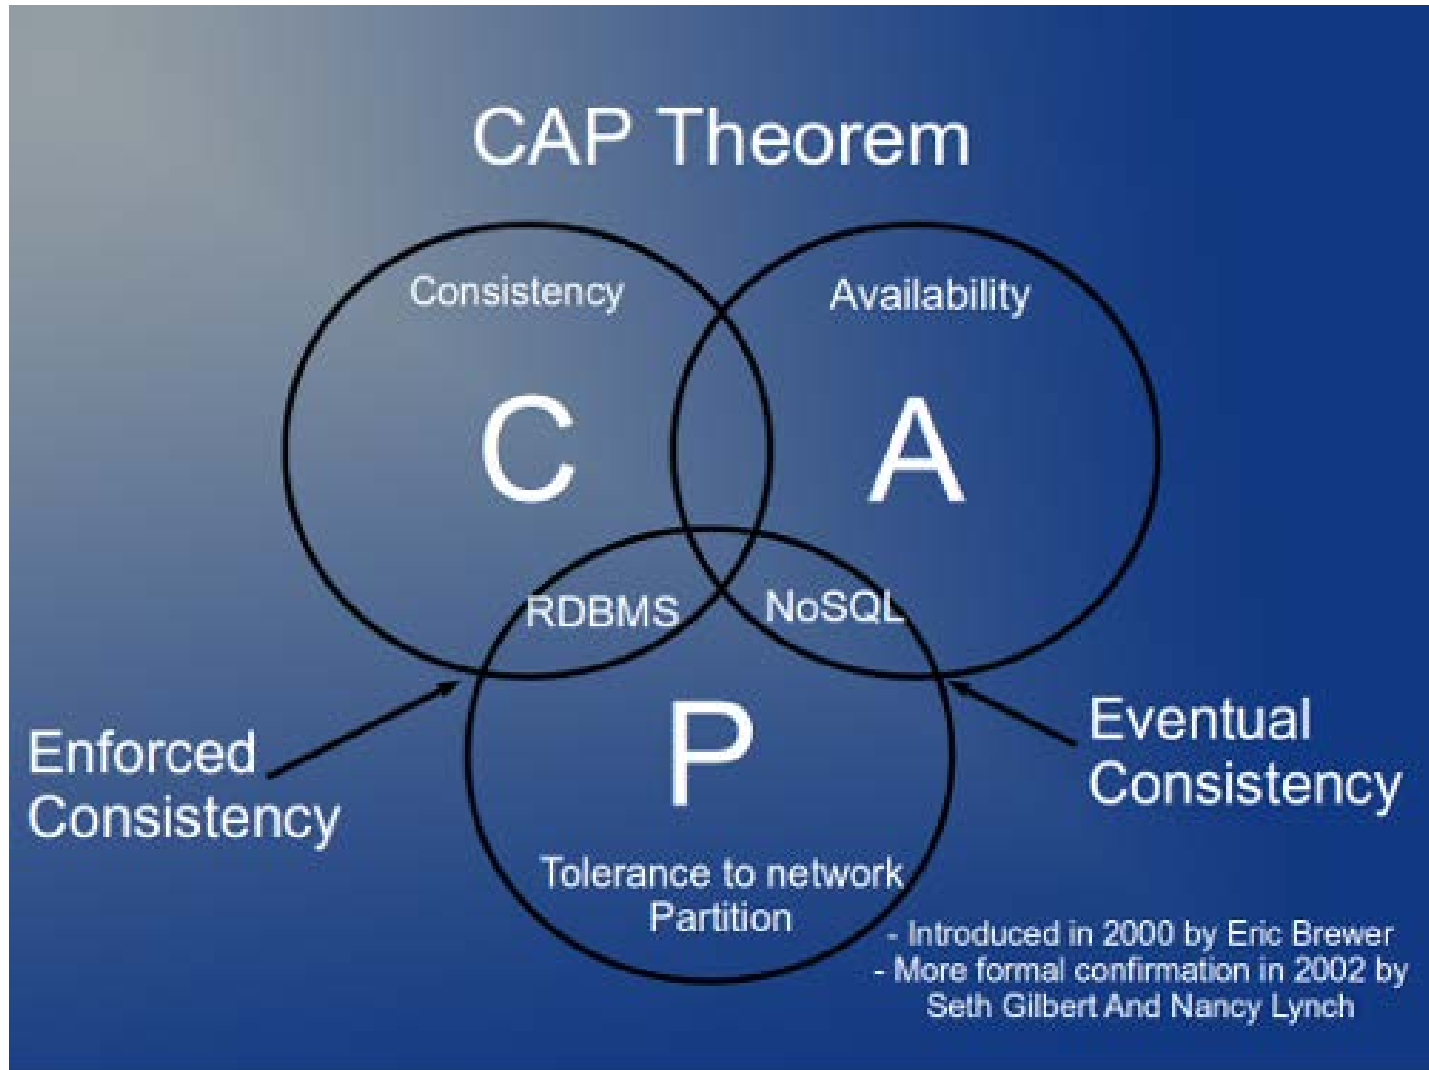
MOTIVATION.

Get off your ass.

Data Lookup pattern challenge



Data Consistency challenge



Data Distribution challenge

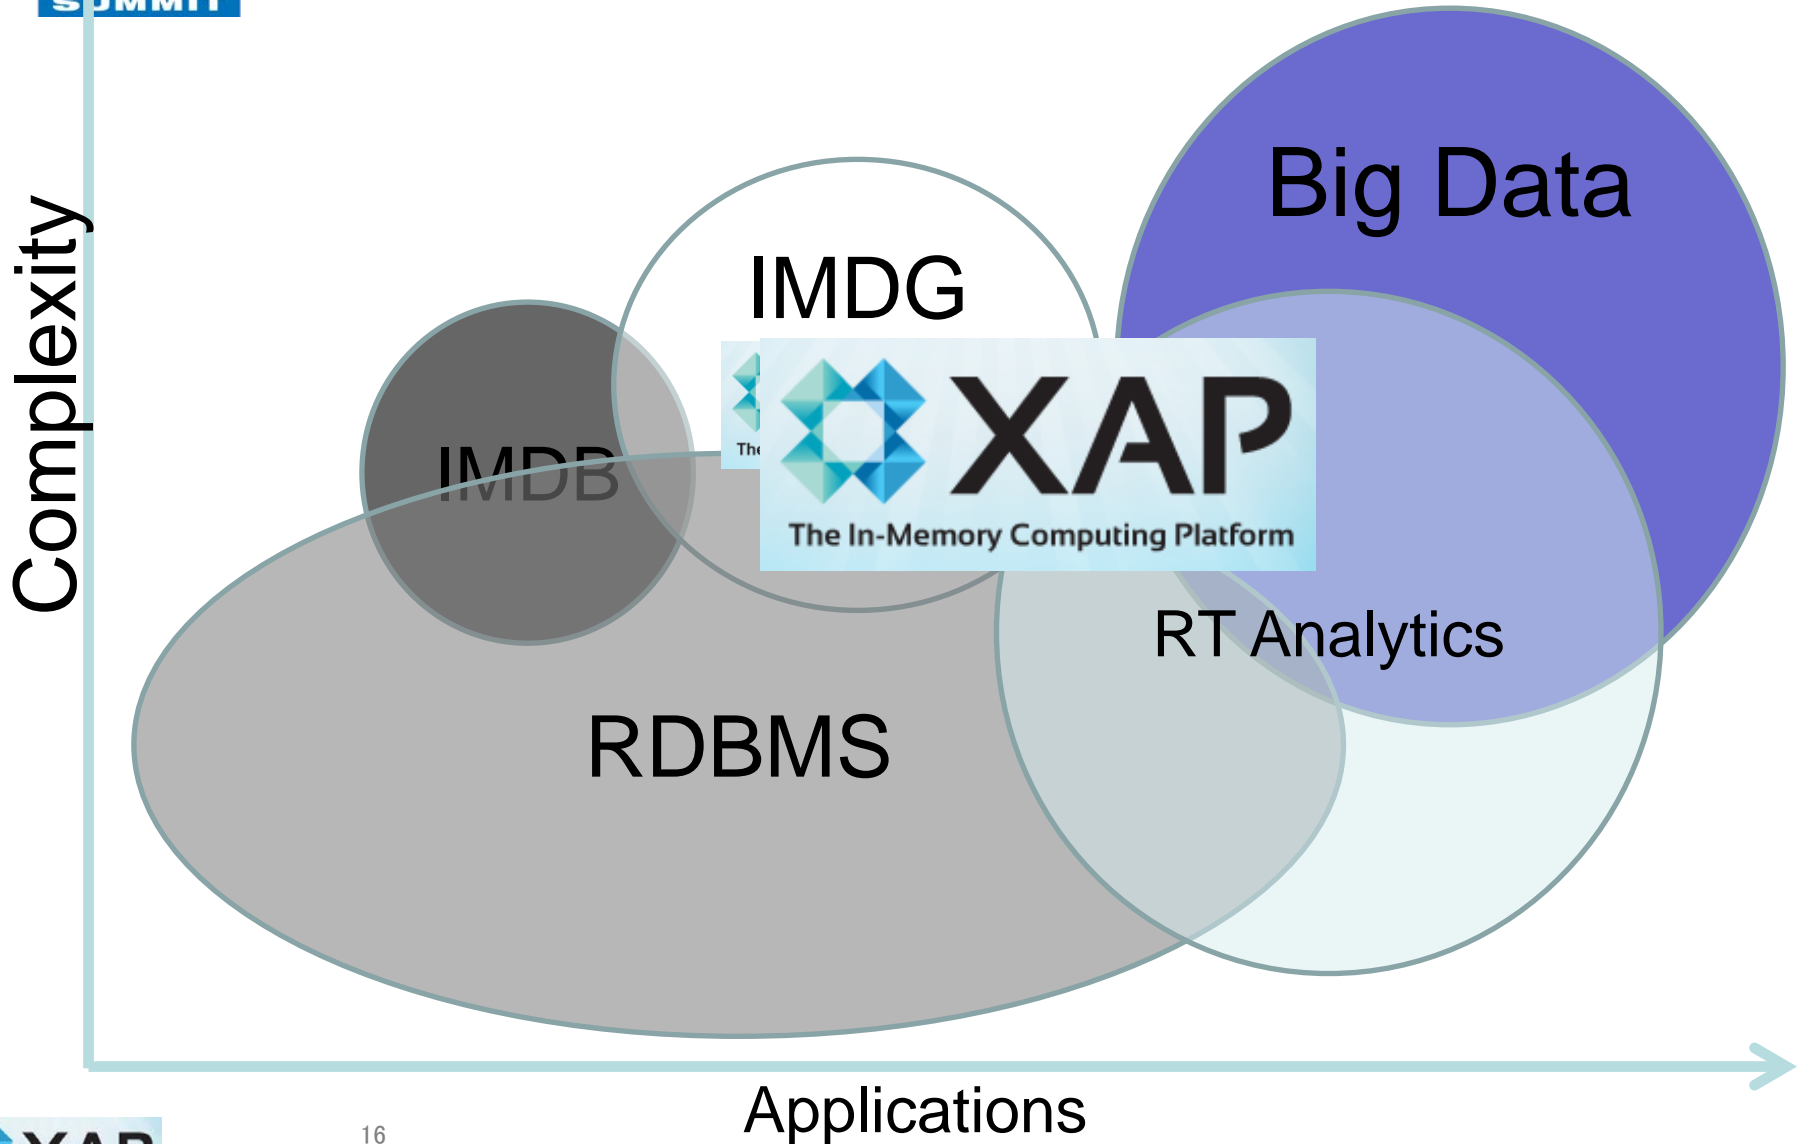


What is XAP MemoryXtend ?

XAP MemoryXtend
manages **large**
amount of data in
grid configuration
using high density
capacity servers.



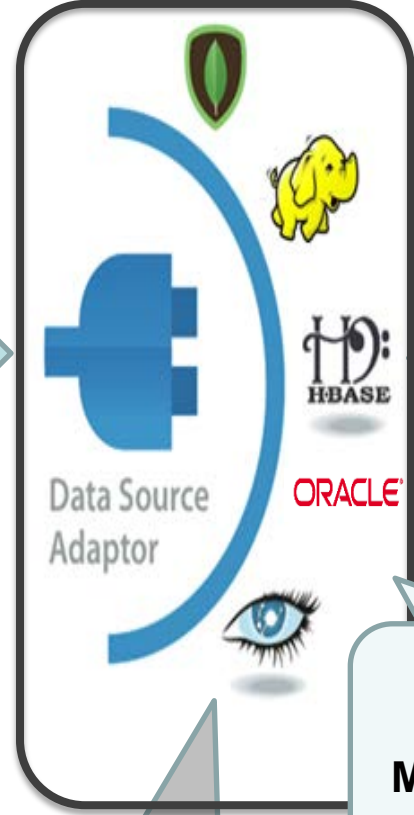
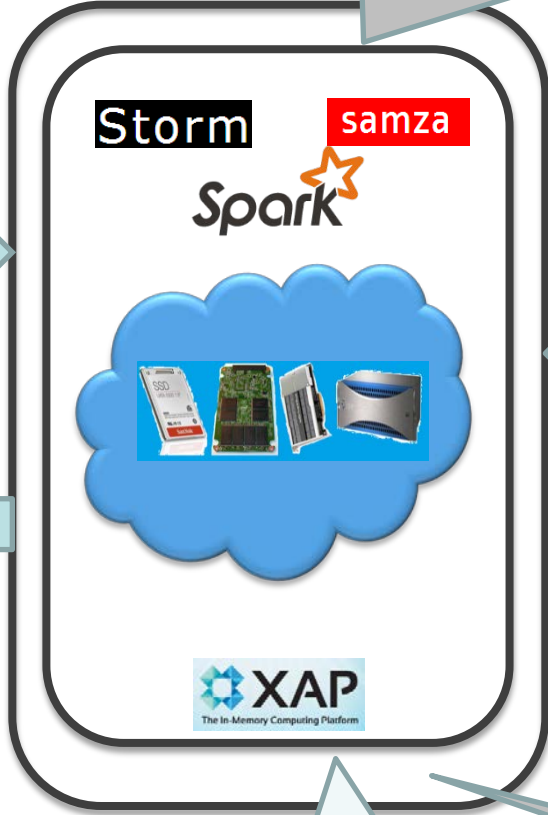
XAP MemoryXtend Position



Typical Usage Scenario - RT Big Data

XAP In-Memory elastic processing -
Counting, Aggregating, Validating ...

Event Stream



batch ANALYTICS



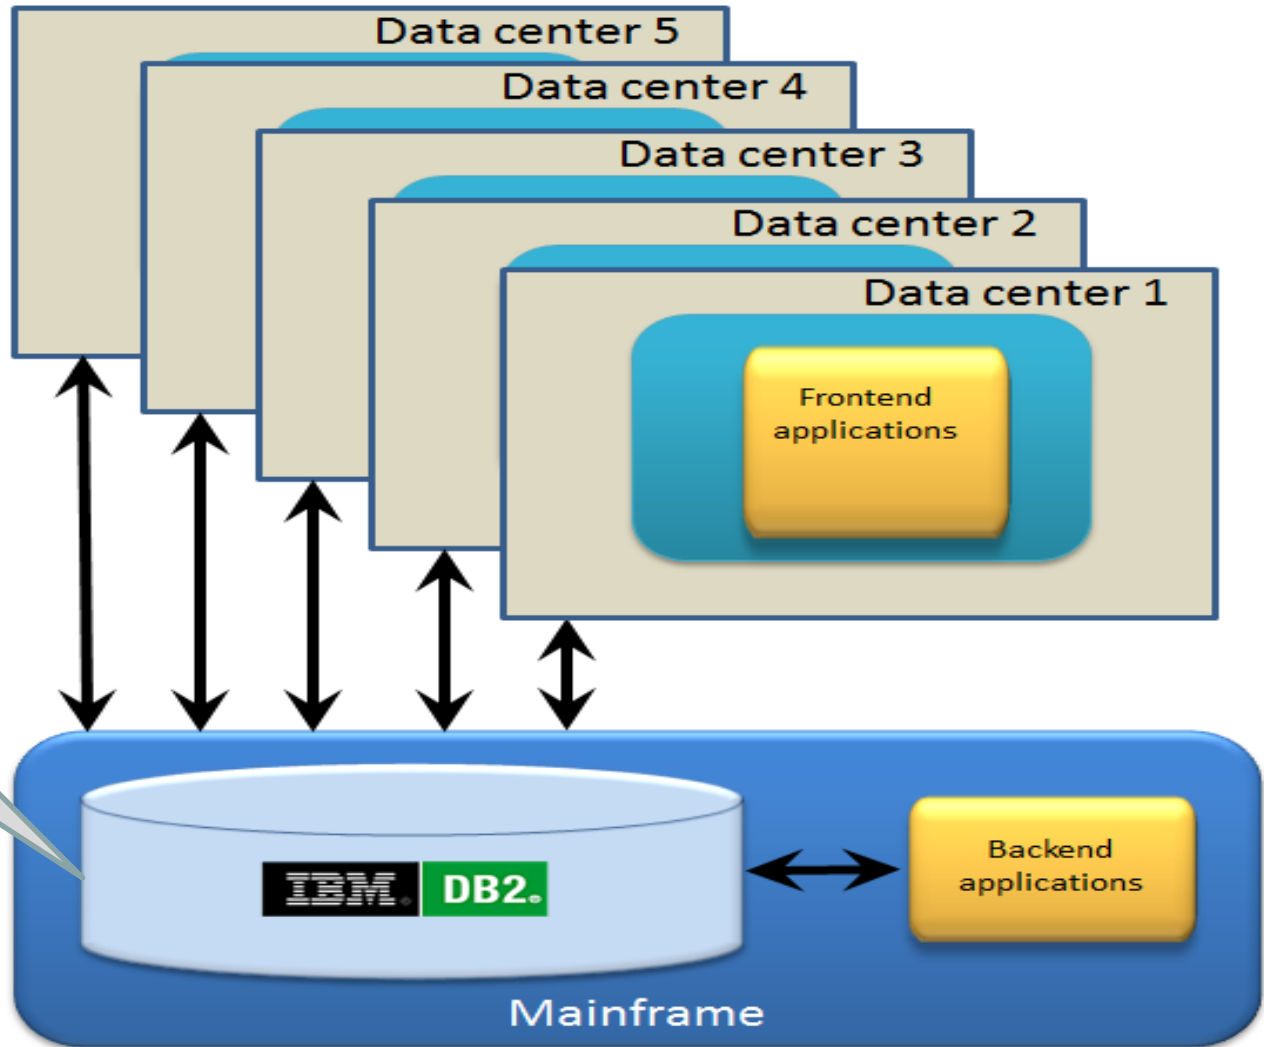
Real time ANALYTICS

User backend environment
Multi PB capacity
- HDD Storage

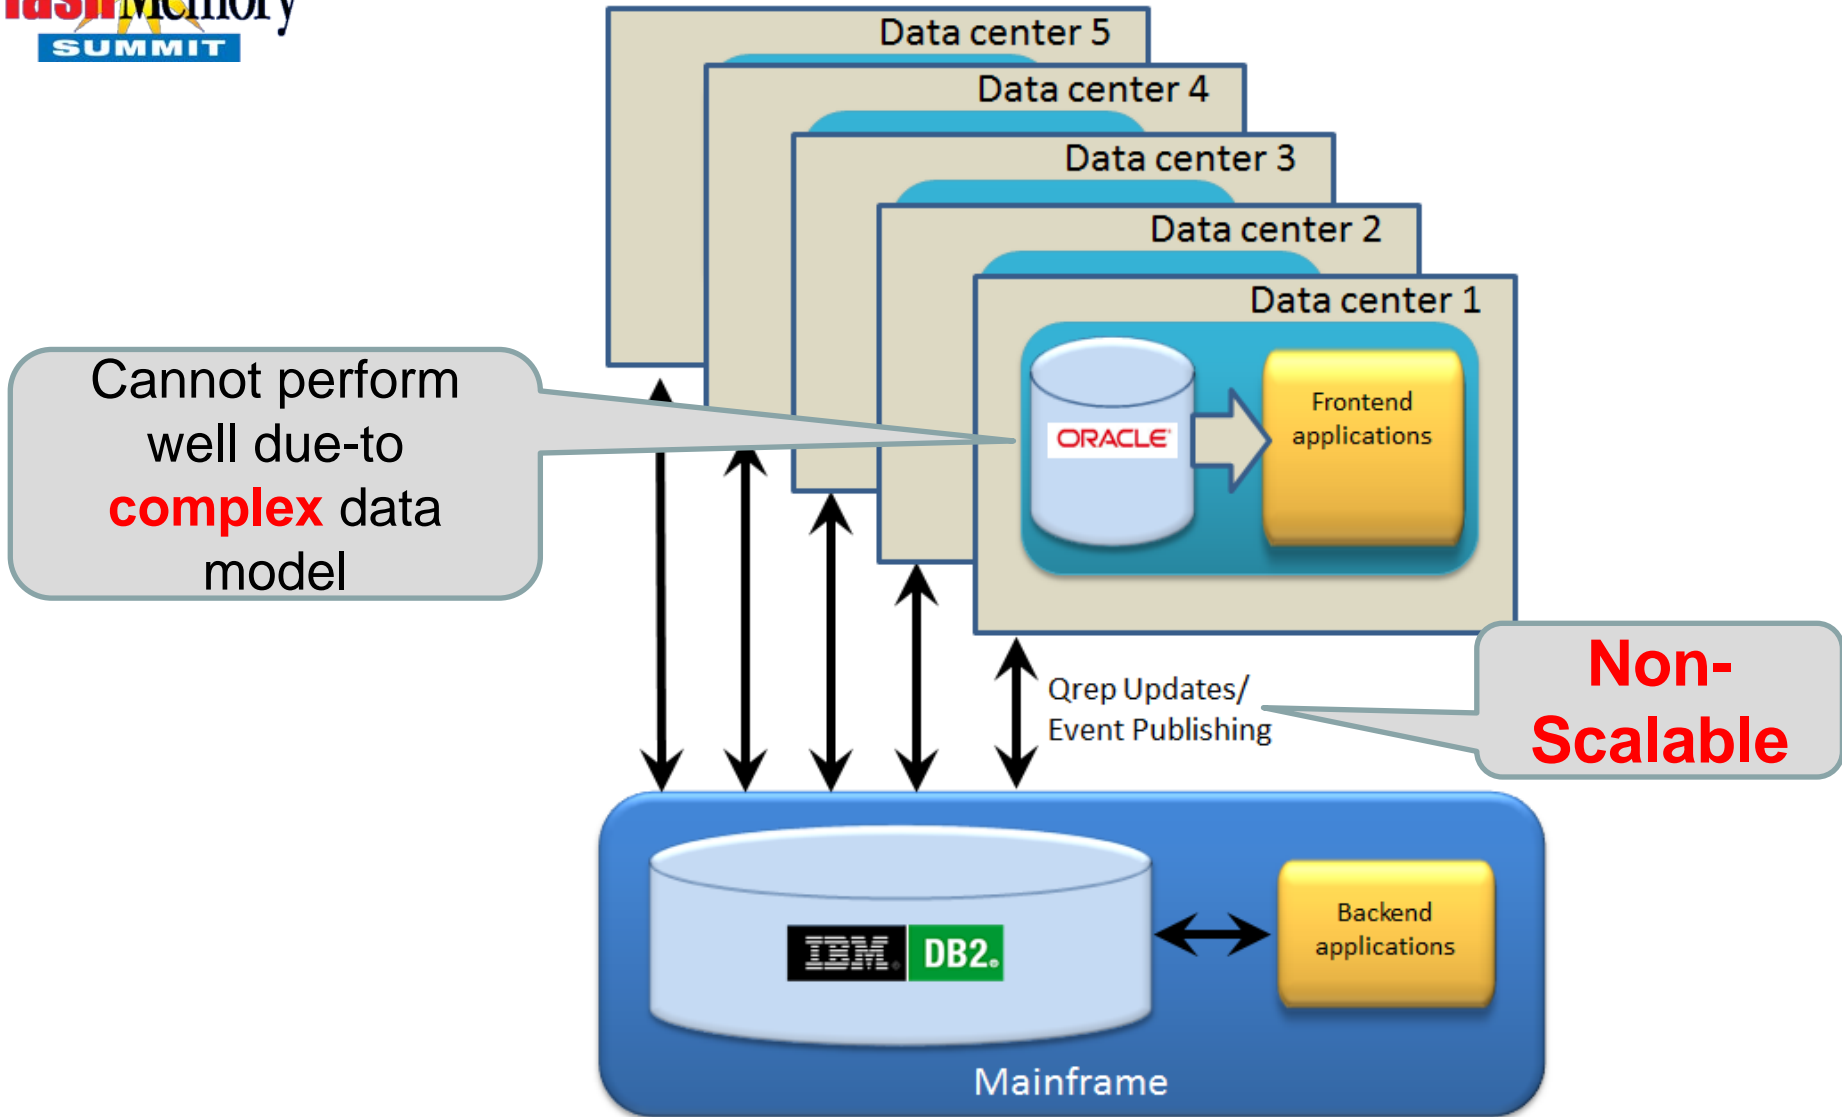
Multi TB capacity -
RAM +SSD Storage

Cludify Elastic management

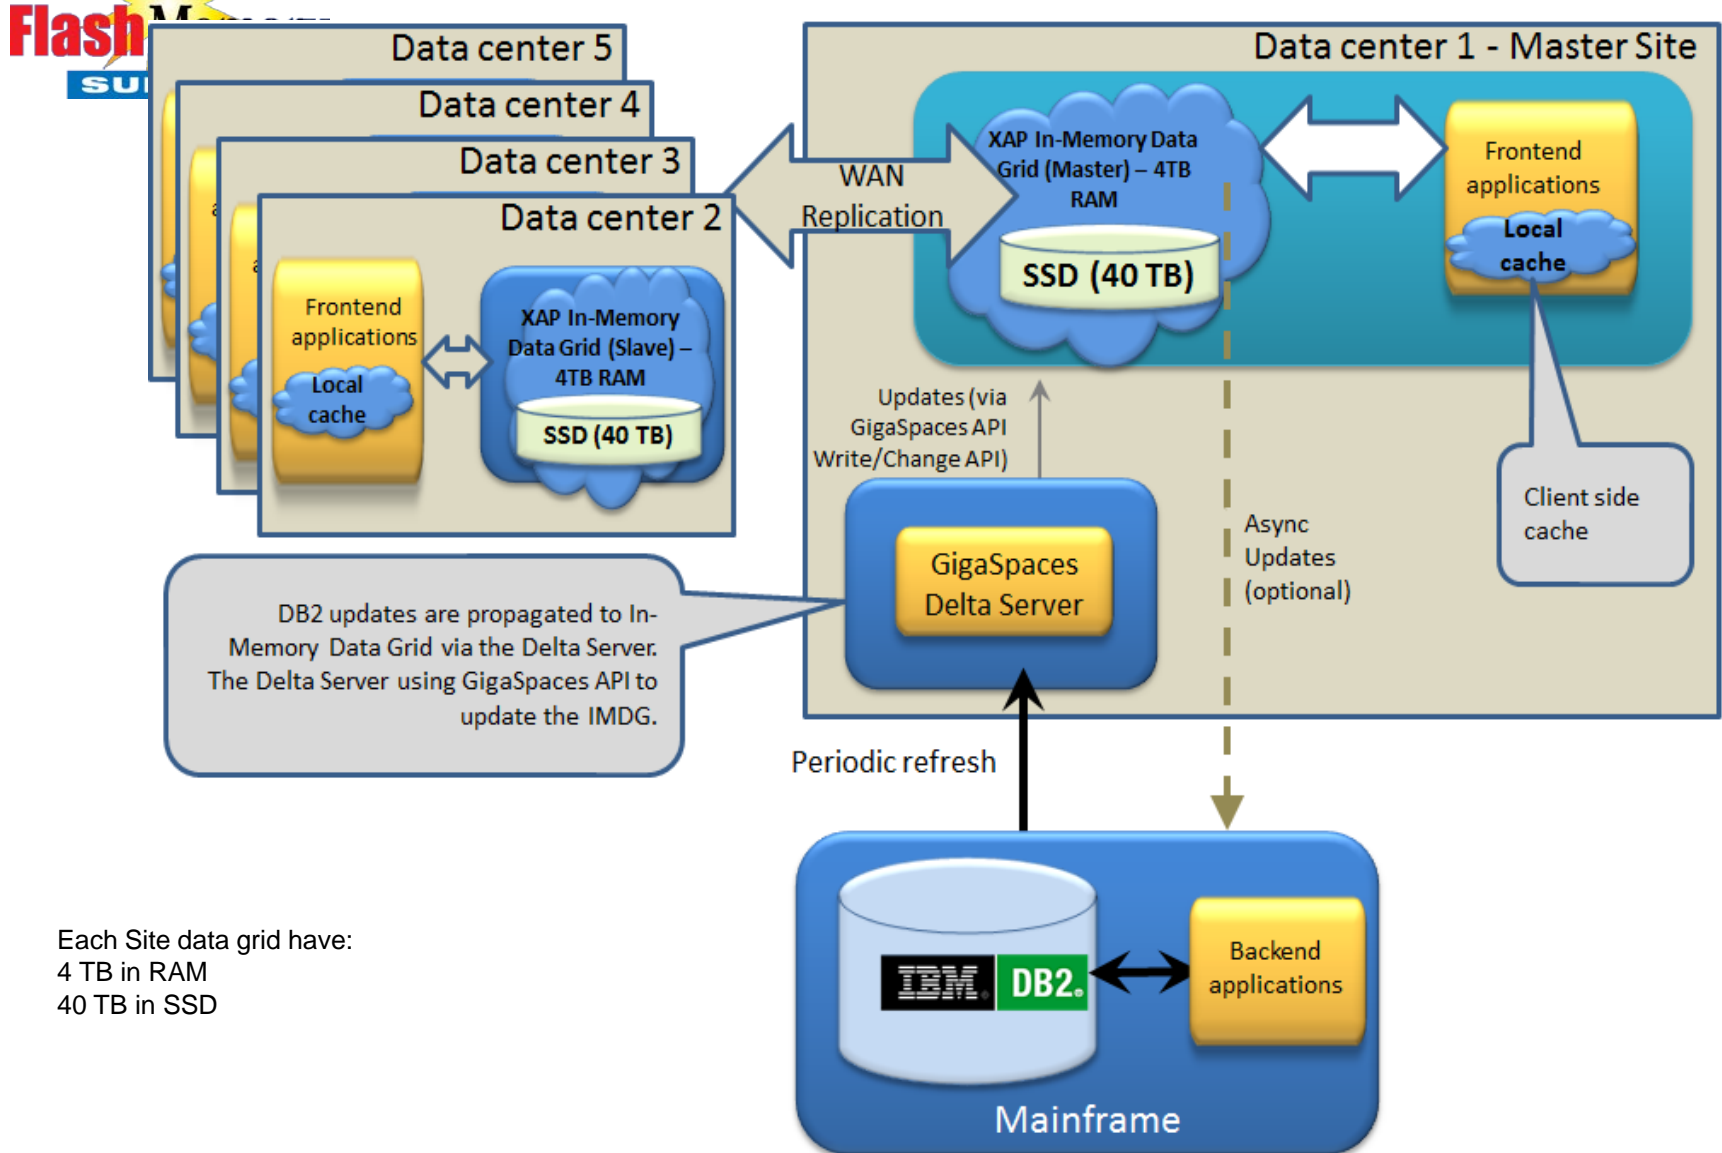
Mainframe Offloading Scenario



Tier 1 Bank eCom Current Architecture



XAP SSD Usage Scenario - Mainframe offload Architecture




Each Site data grid have:
4 TB in RAM
40 TB in SSD

No Oracle DB anywhere , Less Load on the Mainframe !

Performance Benchmark

RAM XAP IMDG vs. SSD XAP IMDG Absolute Write/Read Benchmark Results



SSD is slower than RAM - No surprise!

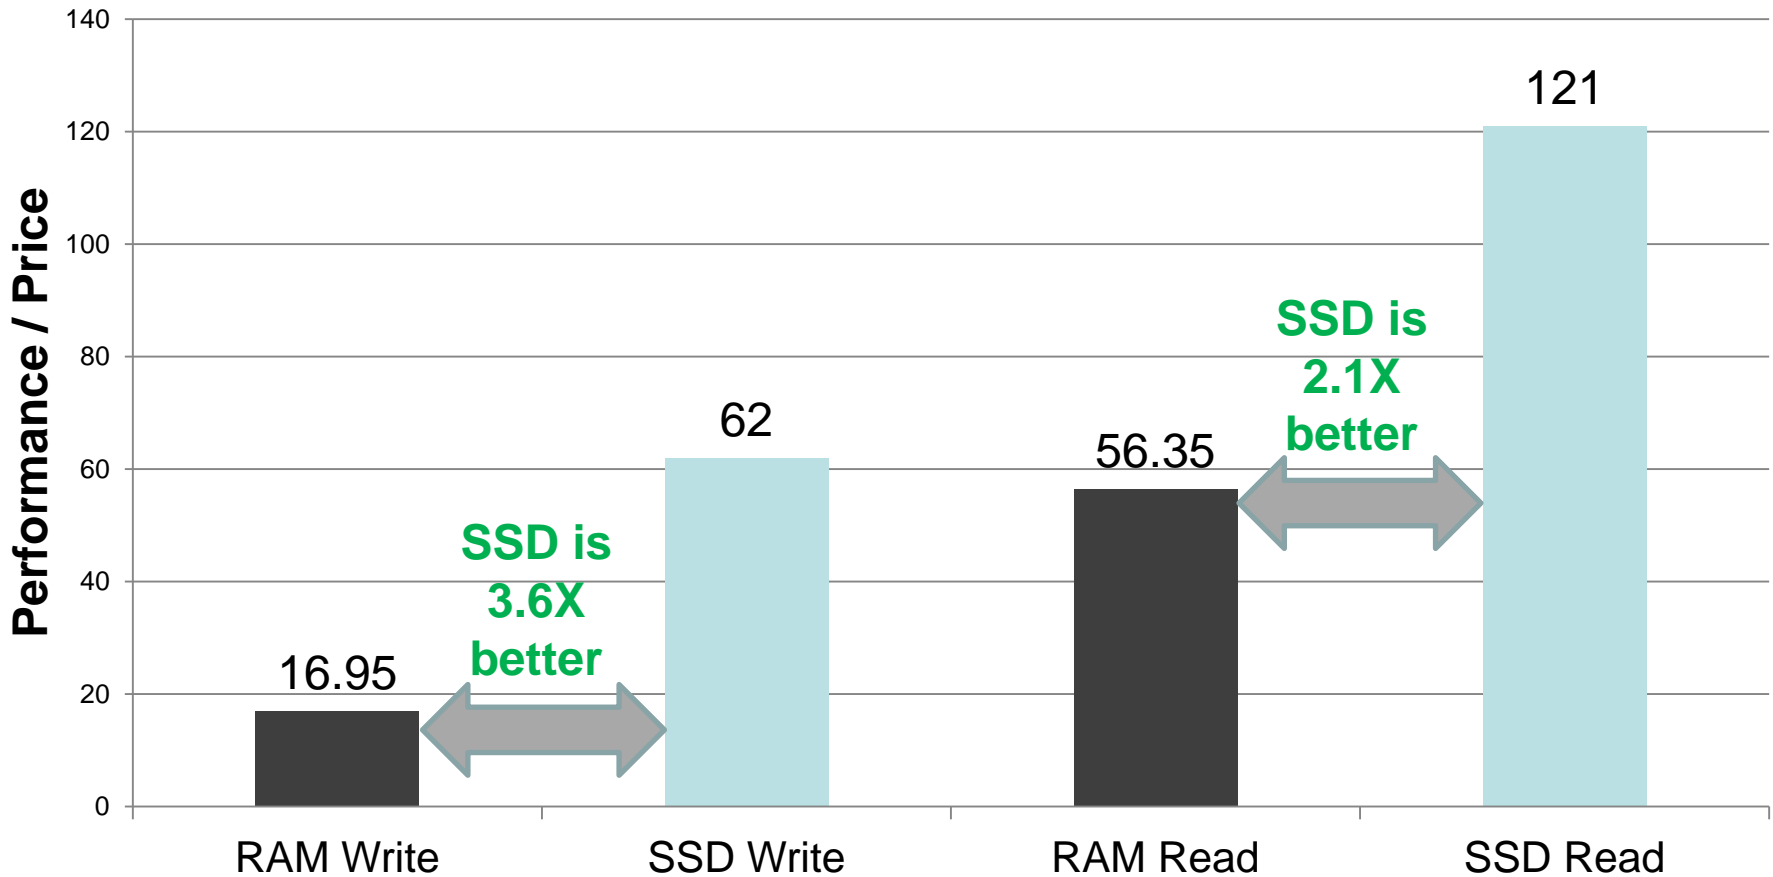
**No database on any HW platform can
deliver such performance !**

* The benchmark was running on HP DL 380 server with 2 sockets 2.8GHz CPU with total 24 cores, 148G DRAM , CentOS 5.8 ,2 FusionIO SLC PCIe cards with software raid 0. The payload is 1KB object size, Single String based key. Uniform read distribution.

* Based on test conducted by SanDisk running on HP DL 380

Price-Performance Benchmark

Price-Performance XAP RAM IMDG vs. XAP SSD IMDG



- We assume 1TB SSD price is \$2K, 1TB RAM price is \$20K.

An Important Difference between the RAM IMDG and SSD IMDG Benchmark...



The RAM data grid benchmark executed with **20GB** total capacity where the SSD data grid benchmark performed with a total of **1TB** data capacity!

SSD Data Grid example...

- [HP DL 580](#) or [Cisco UCS](#) B420/B440 series, comes with built-in support for 2TB RAM and 8 SSD card slots.
- [Intel](#) , [Sandisk](#) and FusionIO or Tegile, offer up to 3TB SSD per card.
- You can have upto 24 TB SSD capacity per server.

**4 servers delivers 8 TB data in
RAM and 96 TB SSD!**

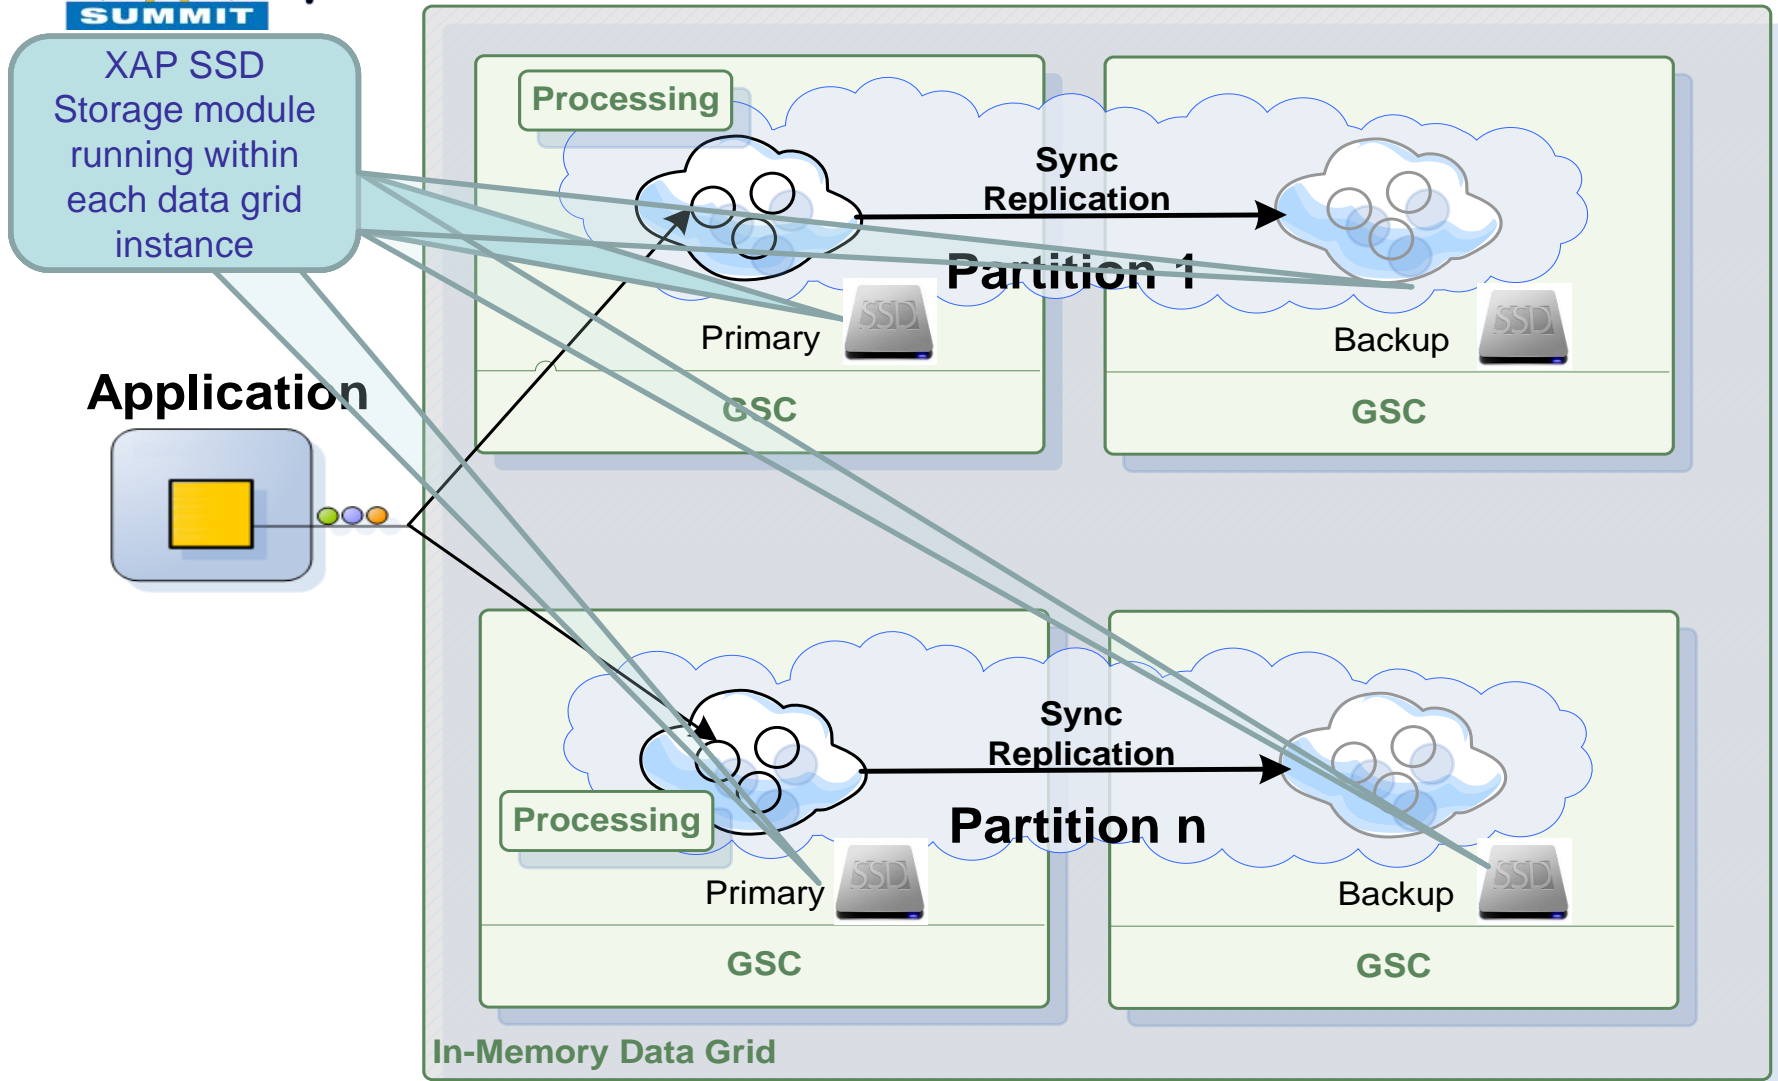
1:10 ratio between RAM:SSD



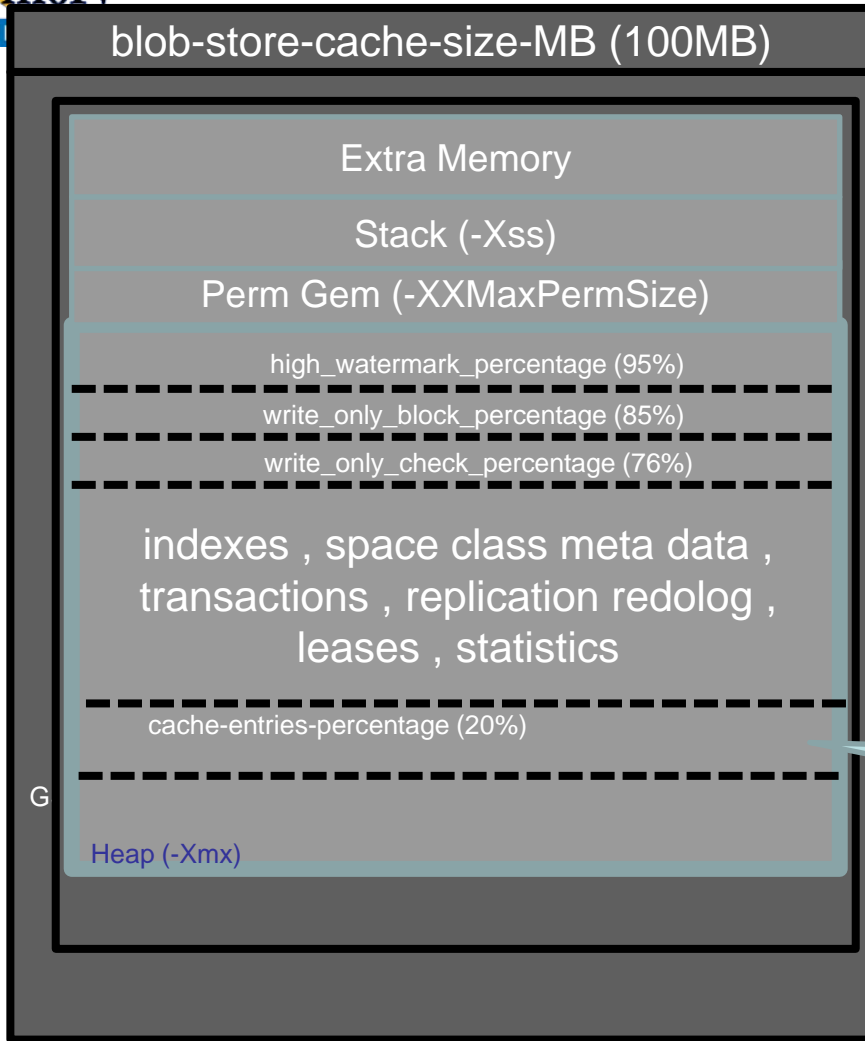
Data Grid Intelligent Application Data Life Cycle Management

- Aging (TTL/Lease)
- Least recently used (LRU)
- First in first out (FIFO)
- Customized (API)

XAP typical Deployment with SSD



XAP Internal Architecture

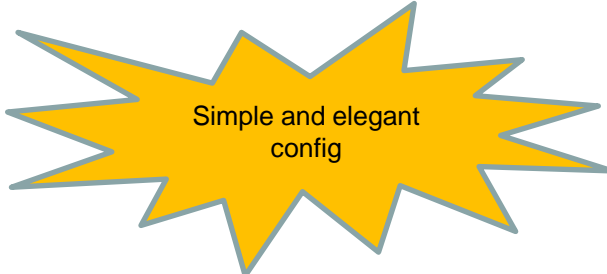


Off-heap in-process serialized LRU cache



On-heap de-serialized LRU cache

Configuration



Simple and elegant
config


```
<blob-store:sandisk-blob-store id="sandisk"  
  blob-store-capacity-GB="50"  
  blob-store-cache-size-MB="20"  
  devices="/dev/sdc1,/dev/sdc2,/dev/sdc3,/dev/sdc4"  
  volume-dir="/data${clusterInfo.runningNumber}"  
  durability-level="PERIODIC"  
  blob-store-reformat="true">  
</blob-store:sandisk-blob-store>
```



Device
mapping

```
<os-core:space id="space" url="/./myDataGrid">  
  <os-core:blob-store-data-policy  
    blob-store-handler="sandisk"  
    cache-entries-percentage="1"  
    avg-object-size-KB="1"/>  
</os-core:space>
```

Dramatic Cost Savings



Server Consolidation

60-80% Reduced HW footprint

100X data per node

From 30GB to 3TB!

Summary: Why XAP MemoryXtend?

- GigaSpaces XAP SSD data grid changing the way we look at mass data storage for large scale real-time applications.
- By leveraging SSD technology together with a standard SSD API and GigaSpaces XAP **MemoryXtend** you can enjoy today an elegant solution that comes with a reasonable price tag.





XAP MemoryXtend

Available for
download from

www.gigaspace.com



Thank you!