

# XAP MemoryXtend - Massive Application Storage Capacity for Real-Time Applications

June 2014

**SanDisk**



GIGASPACE S



**XAP**

The In-Memory Computing Platform

# XAP MemoryXtend-White Paper

## Table of Contents

<b>Introduction</b> .....	<b>03</b>
<b>The Data Challenge</b> .....	<b>03</b>
<b>Overcoming the Data Challenge with In-Memory Computing</b> .....	<b>04</b>
XAP Provides .....	04
<b>The Race for Bigger Data and Faster Data Processing</b> .....	<b>05</b>
<b>The Solution — Combining RAM and SSD Technologies Using XAP MemoryXtend</b> .....	<b>06</b>
SanDisk ZetaScale — Optimizing Access to SSD .....	06
How XAP MemoryXtend Works .....	07
<b>Weighing Performance against Price of RAM and SSD Data Grids</b> .....	<b>09</b>
<b>Analyzing Benchmark Results</b> .....	<b>11</b>
<b>Bottom Line — Which Mode Is Preferable?</b> .....	<b>12</b>
<b>Conclusion</b> .....	<b>12</b>
<b>Opportunities to Evaluate XAP MemoryXtend</b> .....	<b>13</b>
<b>About GigaSpaces XAP</b> .....	<b>14</b>

# Introduction

With growing market demands for faster processing of constantly growing datasets, more and more companies turn to in-memory computing in order to get real-time processing of their data and to obtain application scalability and high availability. In-memory computing is RAM based, and since RAM is costly, in-memory computing entails costs that are prohibitive to many organizations.

This white paper reviews a potential solution for this issue with the newly launched XAP MemoryXtend. XAP MemoryXtend is a combination of GigaSpaces' XAP in-memory computing platform and SanDisk's ZetaScale SSD library, alongside any SSD. This solution enables a hybrid architecture which intelligently stores data on both RAM and SSD, allowing its users to enjoy both high speed processing and cheaper storage capacity. The paper presents a benchmark which weighs XAP MemoryXtend's price against performance and shows that it is indeed the optimal solution in this context.

## The Data Challenge

With more enterprises providing online business services as a main channel of retail and service for their customers, the enterprise application landscape is shifting. Furthermore, a larger percentage of these users are using mobile and tablet devices, which means more concurrent users at any given time. As a result, a growing number of applications need to process hundreds of thousands requests per second. Some of these applications need to provide sub-seconds user experience in order to remain competitive.

Traditional architecture which relies on multi-tier design with a relational database cannot cope with such requirements. As such, there has been a new generation of enterprise middleware products and solutions recently developed, amongst which we find in-memory computing platforms. One of these products is GigaSpaces' XAP in-memory computing platform.

# Overcoming the Data Challenge with In-Memory Computing

The XAP In-Memory Computing platform eliminates the bottlenecks of relational databases and physical storage from the runtime flow of business transactions. XAP is based upon GigaSpaces' In-Memory Data Grid (IMDG), which provides an essential set of data store features such as transactions, indexes, and query language (SQL-like queries).

XAP handles common functions like messaging, event processing, data access, and transaction processing completely and exclusively in memory, and therefore reduces transaction times, allows quick data processing, and speeds up the application. XAP is also interoperable, thus allowing Java, .Net, and CPP.

## XAP provides:

- **Extreme Performance** – XAP stores terabytes of application data in memory, close to the business logic that needs to access this data. There are no transaction delays from physical I/O, database connection pool, or network bandwidth issues. By avoiding conventional disk drives and minimizing network access when performing common middleware functions, XAP delivers performance as much as one hundred times faster and allows a real-time response.
- **Linear Scalability** – XAP enables partitioning of data into self-contained processing units, and provides mechanisms to elastically deploy application units to handle any load and to dynamically allocate resources for optimized utilization. This provides linear scalability, optimized load balancing, and smart resource utilization.
- **High Availability** – XAP guarantees zero downtime, with hot backup and automatic recovery from failure, as well as multi-dimensional monitoring capabilities to quickly locate operational and functional problems.

# The Race for Bigger Data and Faster Data Processing

More and more enterprises today encounter a need to access ever-growing amounts of data and perform analysis on this data in order to improve their offering to customers. Big Data processing is often done offline by tools like Hadoop, which doesn't fully solve the problem. In order to retain a competitive edge, businesses must be able to gain real-time insights and react in real-time or close to real-time. This creates a need for tools that can process these large amounts of data as fast as possible. In-memory computing platforms are a perfect fit in terms of the speeds at which they can operate and process the data.

As an IMC platform, XAP is able to perform in-memory processing for Big Data and provide real-time insights for the most complex scenarios; however, it relies completely on RAM, which is limited in the capacity of data that it can store. XAP has an upper bound limit of a few terabytes.

This is where solid-state drives (SSDs) come in. Since SSD prices have reduced significantly, they are now far more widespread. SSDs today are no longer limited to consumer devices such as laptops and tablets, but are increasingly used as an infrastructure to enterprise class storage devices in modern data-centers. They are typically used to store relatively transient data, such as transactional data that is inherently smaller and has a short lifespan, alongside highly concurrent access patterns. SSD drives indeed provide a cheaper storage alternative compared to RAM; nevertheless, accessing data with a SSD is typically ten to one hundred times slower than accessing the same data using RAM. Would it not be ideal if there was a way to combine RAM's performance with SSD's capacity?

# The Solution — Combining RAM and SSD Technologies Using XAP MemoryXtend

XAP MemoryXtend leverages XAP's IMC capabilities alongside SanDisk's ZetaScale technology — a unique SSD access library which optimizes the access to SSD drives of any kind and achieves access speeds typically two to ten times faster than ordinary file-based access. By leveraging ZetaScale to extend XAP built in RAM capabilities, XAP MemoryXtend supports datasets of tens and even hundreds of terabytes. It enables storing, processing, and analyzing data at speeds comparable to RAM, with significantly reduced TCO and a decreased server footprint.

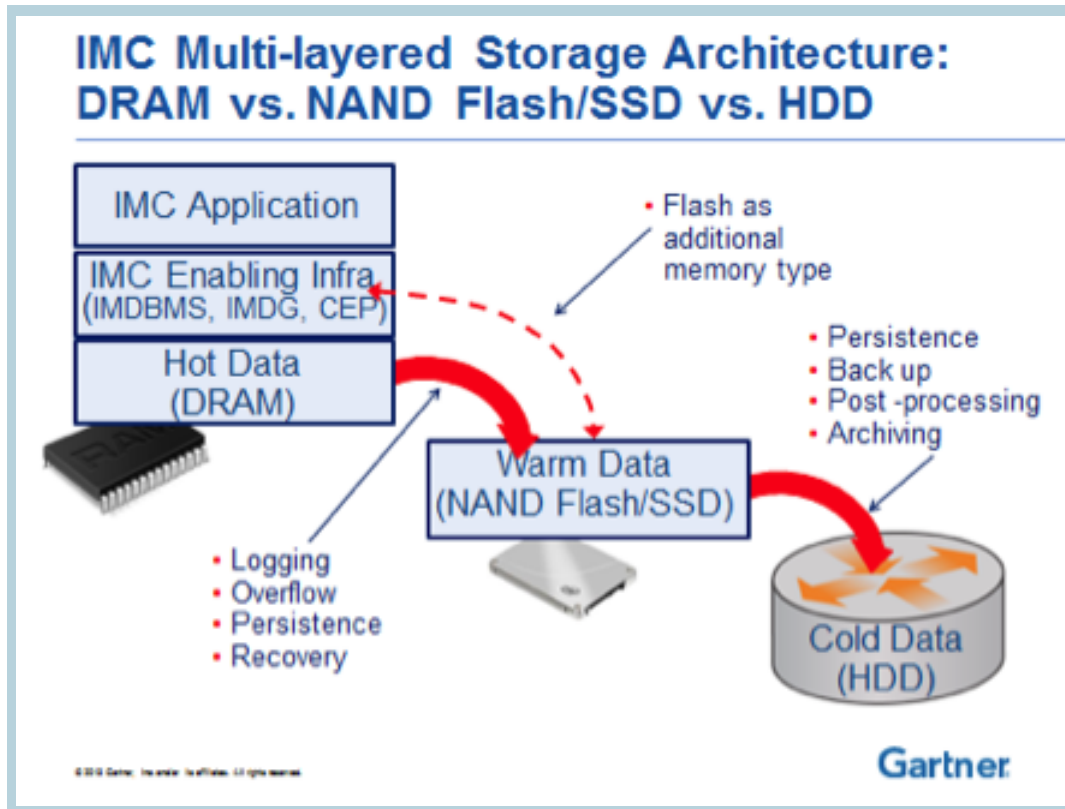
## SanDisk ZetaScale — Optimizing Access to SSD

SanDisk ZetaScale software technology allows applications to utilize the full benefit of flash data storage and parallel data access. When combined with an in-memory computing platform and with any SSD drive, ZetaScale can provide data access speeds close to that of DRAM at a fraction of the cost.

Applications access ZetaScale through an API that can very efficiently fetch, cache, and store individual application objects, or records. ZetaScale selectively utilizes DRAM and flash for data most appropriate for each storage medium, caching frequently accessed objects in DRAM and efficiently buffering and aggregating flash writes to minimize flash wear. ZetaScale exploits the high capacity low-latency persistence, parallelism, and high throughput of flash memory.

The combination of ZetaScale together with XAP IMC creates a distributed data fabric that can store application data across RAM and SSD in a transparent manner. XAP MemoryXtend is an extension to the in-memory data grid that provides the same XAP application-level data abstraction with a variety of data access APIs (Object/SQL, JDBC, JCache, JMS, and NoSQL). The extension leverages ZetaScale's optimized SSD access capabilities. By providing the same scale out capabilities for memory only deployments, XAP MemoryXtend enables users to store and traverse large amounts of data, while maintaining the ability to increase capacity as needed, without any architectural changes or downtime.

## How XAP MemoryXtend Works



As can be seen in the diagram above, the XAP In-Memory Computing layer acts as the front-end to the flash device and handles the transactional data access and stream processing, while the flash device is used as an extension of the RAM device. From the application's perspective, this is all transparent.

There are basically two modes of integration between the RAM and SSD:

**LRU Mode** – Using the RAM as a caching layer to the SSD, the RAM holds the “hot” (the most recently used data), and the SSD holds the entire data set. LRU mode is optimized for maximum capacity; however, it has a limited query interface of simple key/value data access.

**Fast Index Mode** – In this mode, all the data is stored on the SSD drive, but its indexes are stored in RAM. The Fast Index mode supports complex queries, including range queries and aggregate functions; however, it is limited to the size of indexes that can be held in RAM.

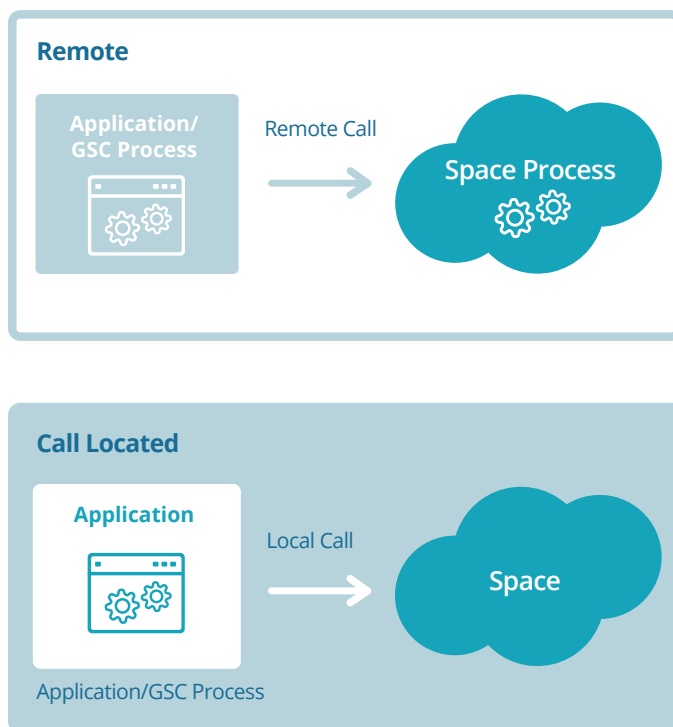
In both cases, the access to SSD is done directly, using ZetaScale's Key/Value interface and not through the ordinary file system APIs. Since both the RAM and the SSD are fast, data is written synchronously to the SSD and this way any potential inconsistency is avoided. It is also quite common to synchronize the data from the IMC and flash devices into an external data store that runs on a traditional hard drive. This process is done asynchronously and in batches to minimize the performance overhead.

Optimizing access to the SSD and using it as part of a hybrid architecture brings server consolidation and therefore cuts costs. However, one issue still remains: in terms of performance, RAM is faster than SSD. The following pages discuss the tradeoffs of price versus performance in further detail, backed by benchmark results which prove that this type of hybrid architecture is indeed the optimal choice in this context.



# Weighing Performance against Price of RAM and SSD Data Grids

A data grid that is running purely in RAM on a commodity server can perform around one million write/read operations per second per node (without replication), with a 256-byte payload. A client-side cache can actually deliver even faster read performance. This type of data grid node should run on a commodity server with the latest multi-core CPU, as it must support a highly concurrent environment with many clients/threads accessing the data. If, for example, a capacity of ten million writes per second is required, one can simply run a clustered data grid across ten nodes. It is possible to scale linearly.

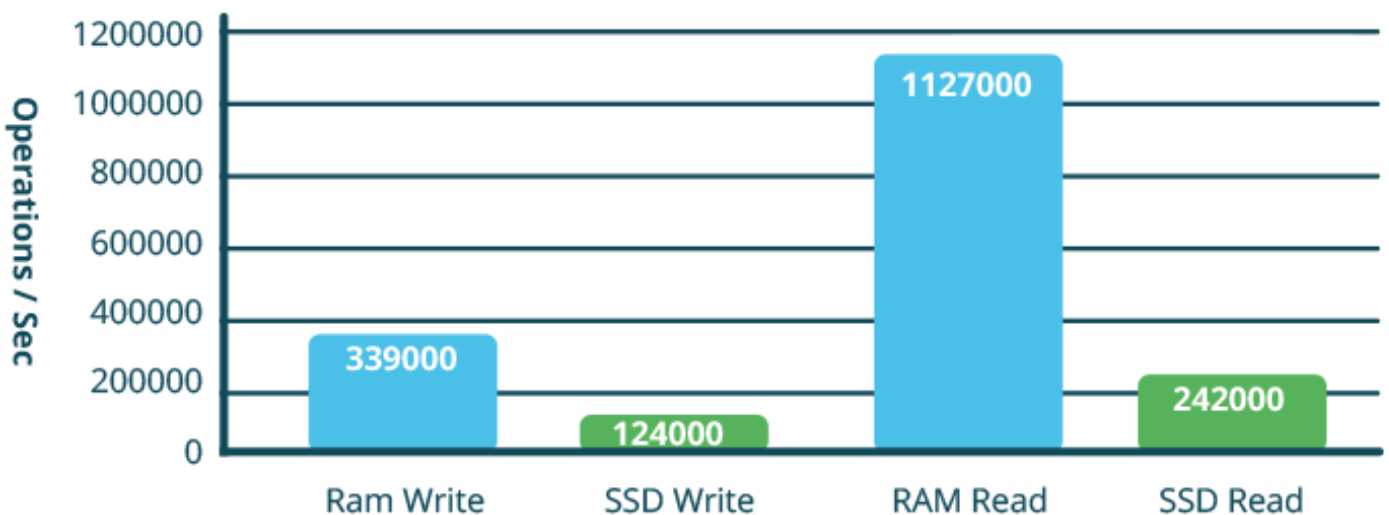


Data grid read/write operations can be executed by remote clients or collocated clients (similar to database-stored procedures or triggers). With collocated clients, the operations will enjoy ultra-low latency, as there is no serialization or network activity performed (approximately ten microsecond latency). Remote operations will have around one millisecond latency when accessing the data grid (depending on the payload size), which is still reasonably fast.

This benchmark was based upon the [Yahoo cloud benchmark \(YCSB\)](#). It compared a pure RAM data grid with a data grid using the XAP MemoryXtend technology when performing write and read operations in a collocated configuration. With the MemoryXtend-enabled data grid, read operations were performed without warming the data grid. This means that all read operations consumed their data from the SSD (100 percent cache miss).

GigaSpaces XAP delivered 1.2 million read TPS in pure RAM mode and 242K read TPS in SSD mode, and 340K write TPS in RAM mode and 120K write TPS in SSD mode. This indeed proves that RAM is faster than traditional SSD, but also that using this configuration the orders of magnitude are smaller than typically imagined.

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



# Analyzing Benchmark Results

When analyzing the results, it becomes clear that when it comes to write operations, XAP running in SSD mode delivers 3.6 times better price-performance than XAP running in RAM-only mode. For read operations, the SSD mode delivers 2.14 times better price-performance when compared to RAM mode. The conclusion, therefore, is that the MemoryXtend technology is actually very cost effective:

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



With this benchmark, write operations against the SSD data grid were performed using a new storage interface provided by XAP, where the write operation is fully acknowledged when both the data grid and the SSD have performed the transaction.

The benchmark was running on a HP DL 380 server with two sockets (2.8GHz CPU) with twenty-four cores, 148G DRAM, CentOS 5.8, and two FusionIO SLC PCIe cards with software raid zero. The payload is 1KB object size, single-string-based key, and uniform read distribution.

# Bottom Line — Which Mode Is Preferable?

There is one important difference between the RAM and SSD data grid benchmark results: the SSD data grid benchmark was performed with significantly more data. The RAM data grid benchmark executed with a total capacity of twenty gigabytes, whereas the SSD data grid benchmark was performed with a total capacity of one terabyte — a fifty-fold capacity difference.

The result of this fast and optimized application level SSD interface is that the SSD data grid can deliver a comparable performance, with a significantly bigger storage capacity per node. Data updates and data removal operations that have traditionally impacted performance are also less painful here as the ZetaScale technology continuously reorganizes the data on the SSD to avoid fragmentation.

Adding into the mix the latest advances in JVM garbage collection and the smooth operation brought about by the [garbage first policy](#) (G1), the end result is a distributed elastic data grid with virtually no pauses managing vast amounts of data per node.

## Conclusion

Big Data mission critical applications require large-scale data grids with fast response times that come at an affordable price. GigaSpaces' XAP MemoryXtend provides just that. By leveraging SSD technology together with SanDisk's ZetaScale, XAP MemoryXtend provides an elegant and cost-effective solution. Enterprise-grade applications can leverage a data fabric that delivers real-time response times, and supports an application-native object oriented data model, full SQL support, and a common application programming API. Enterprise systems are no longer limited to the use of RDBMS or raw files that are centralized and utilizing old spinning disk technology; rather, they can leverage a distributed, mature, elastic data grid which delivers high performance while intelligently keeping data on both RAM and SSD.

# Opportunities to Evaluate XAP MemoryXtend:

Download a trial version of XAP MemoryXtend at no cost at <http://www.gigaspace.com/xap-memoryxtend-flash-performance-big-data>

Try XAP MemoryXtend now using our lab at no cost — book your lab time at <http://www.gigaspace.com/xap-memoryxtend-flash-performance-big-data>

# About GigaSpaces XAP

GigaSpaces XAP is an in-memory computing application platform designed for real-time Big Data analytics applications. GigaSpaces XAP is used by top Fortune 500 companies from all industries, including financial services, ecommerce, telecommunications, and travel. Read more about the XAP In-Memory Computing platform at <http://www.gigaspaces.com/xap-in-memory-computing-event-processing/Meet-XAP>

## References

<http://arstechnica.com/gadgets/2012/12/ssd-prices-are-low-and-theyll-get-lower/>  
[http://en.wikipedia.org/wiki/Solid\\_state\\_drive](http://en.wikipedia.org/wiki/Solid_state_drive)

**Thank You**

[www.gigaspaces.com](http://www.gigaspaces.com)