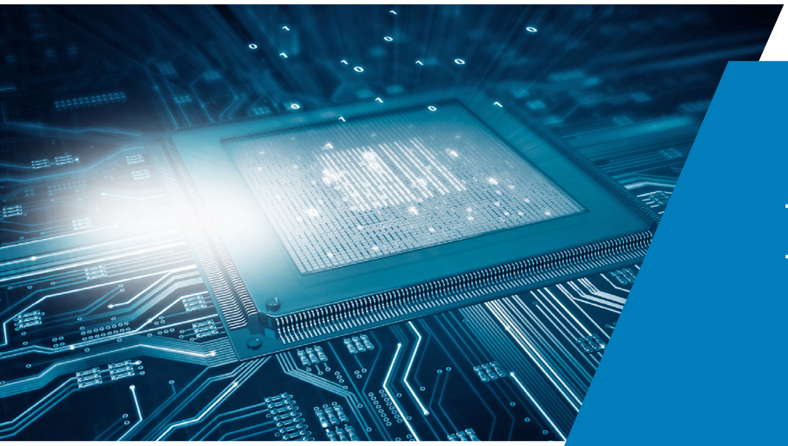

In-Memory Computing : Premiere

Overview

Modern business environments have been evolving rapidly, pushing IT needs and expectations to new highs. Businesses are not asking for performance alone, but extreme performance and, that too, at scale. Traditional application architectures and infrastructures, storage and compute infrastructures, and IT delivery models have all been evolving rapidly in response to these business demands.

A key, recent vector in this evolution is the rapid rise of In-Memory Computing (IMC) as a viable technology alternative to traditional performance- and scale-enabling approaches.

This white paper examines the emerging landscape of in-memory computing, the key solution stacks, and leading IMC technology vendors. It also assesses the readiness of IMC for enterprise grade adoption.

Background

In the past, in-memory computing focused mainly on distributed caches, in-memory data grids, and platforms that enabled space-based architecture. In last couple of years, however, the IMC space has undergone dramatic changes. It has gained immense traction from research community and industry.

Leading software vendors, including platform providers and application software providers, are now deeply entrenched within IMC in one way or the other. As an example, SAP has come up with in-memory analytical database called HANA; TIBCO has ActiveSpaces which is an in-memory data grid; Oracle offers in-memory RDBMS called TimesTen database and Distributed Cache called Coherence; Gigaspaces provides extreme transaction processing platform called XAP; and, recently Microsoft has announced Hekaton, a project code name for SQL Server in-memory OLTP. This whitepaper highlights recent developments and trends in the IMC space.

LEADING VENDORS: IMC TECHNOLOGIES

Let us start our assessment with a quick run through of the leading vendors in the IMC segment and their respective focus areas.

This assessment is based on the publicly released Gartner Cool vendor list for IMC.

#.	Vendor	Year	Description	Product Classification	Key Competitors
1.	Diablo Technologies*	2014	Diablo provides memory channel storage that gives OEM the ability to design low latency SSD by utilizing DDR3 dual in-line memory module.	DIMM Semi-conductor	IBM
2.	Gridgain*	2014	At core, offers IMC platform, IMDG, IMDB, In-Memory Hadoop Accelerator.	IMC, IMDG, IMDB	Gigaspaces, Scaleout State Server, GEMFire, etc.
3.	MemSQL*	2014	Offers in-memory database which belongs to the hybrid transactional and analytical processing database class.	HTAP	SAP HANA
4.	Relex*	2014	Offers an in-memory-enabled, integrated SCP platform for the retail and wholesale industries.	IMDB	JDA
5.	CloudTran*	2013	Offers transaction management product and is used to coordinate ACID properties for data stored in IMDG such as Oracle Coherence and Gigaspaces.	In-memory transaction processing	Kabira Technologies
6.	Garantia Data (now Redis Lab)*	2013	Garantia data has renamed itself to Redis Lab. It offers Redis Cloud and Memcached Cloud. They are in-memory key value data store or NoSQL database.	In-memory NoSQL database	Aerospike
7.	Jedox*	2013	Offers self-service BI and analytical solutions based on in-memory GPU technology.	In-memory OLAP engine	Quartet FS
8.	Quartet FS*	2013	Offers a product called ActivePilot which performs in-memory analytics.	In-memory OLAP engine	Jedox
9.	Panopticon (acquired by DataWatch Corporation)*	2013	Offers visual data discovery and analysis software for real-time, CEP, and historical time series data. It incorporates unique StreamCube™ OLAP data model which provides rich in-memory OLAP functionality for fast analysis for static, streaming, and time series data sets.	In-memory OLAP engine	Tableu Software, Tibco SpotFire
10.	Kabira Technologies*	2013	Offers in-memory transaction processing platform.	XTP: Extreme transaction processing platform	CloudTran
11.	Hazelcast**	2012	Offers in-memory data grid.	IMDG	JBoss Infinispan

#.	Vendor	Year	Description	Product Classification	Key Competitors
12.	PrismTech**	2008	Offers platform for distributed data sharing.	Kind of in-memory messaging platform	
13.	Terracotta***	2006	Offers distributed caching product.	Distributed cache	Oracle Coherence

List: IMC** ****List: Application & Integration Platform** ** List: Platform Middleware**

ANALYSIS OF IN-MEMORY TECHNOLOGIES

In the past, IMC products were limited to point solutions largely focused on distributed cache and in-memory data grid systems. The dramatic change witnessed in the recent years is the reincarnation of many traditional products by deeply reengineering them to leverage the benefits of ultra-high performance DRAM as the primary data store.

For example, traditional relational databases were disk based for delivering ACID (Atomic, Consistency, Isolation and Durability) properties of transactions. Today, many vendors, such as Oracle and VoltDB, are offering in-memory relational databases (IMRDB) with ACID properties. On the similar lines, many traditional disk-oriented products are now going the in-memory route to enable applications to deliver high performance.

Let us analyze some of these product segments.

1. 1. Relational online transaction processing (OLTP) databases

In order to provide transactional capabilities with ACID properties, traditional relational OLTP databases were always disk based. A few other reasons that traditional OLTP databases were disk based are: DRAMs were expensive, computers had very little memory, and memory-based technologies were not adequately mature .

For preserving data integrity, traditional databases add many processing overheads such as index management, WAL (write-ahead logging), locking, and buffer management, all of which limit performance and prevent these systems from scaling horizontally to handle large data volumes and workloads. IMRDB (in-memory relational database) such as VoltDB, on the other hand, essentially

> **From fragile and expensive point solutions for niche use cases to robust and highly affordable alternates, IMC stacks have evolved radically. Simpler architectures, extensive scalability and massive and predictable performance are direct benefits of IMC stacks.**

removes these processing overheads by keeping data completely in memory and employing memory-based architecture for database. For example, keeping data completely in memory eliminates the need for buffer and index management. Further, implementing single-threaded partitions helps IMRDB to eliminate expensive locking. Finally, replication and command logging allows IMRDB to achieve high availability and durability despite being memory oriented.

However, IMC adoption is no more an exclusive trend among new DB vendors. Established disk-based OLTP database vendors have already started appreciating and building extensive IMC mix into their solutions. Oracle (TimesTen DB), IBM DB2 BLU, Microsoft (Hekaton DB), IBM SolidDB (now a UNICOM product), to name a few, are already in the fray.

2. NoSQL databases

NoSQL databases emerged in response to the business needs of supporting massive data scale and load that could not be satisfied with traditional disk-based OLTP databases owing to their inherent limitations.

The relational OLTP database is a scale-up technology. They are difficult to scale-out because, by design, they have centralized share-everything architecture. NoSQL databases are designed based on distributed share-nothing principles, which allow seamless horizontal scaling. This allows NoSQL databases to handle millions of concurrent users and several hundreds of terabytes of data. However, typically, these NoSQL databases tradeoff ACID qualities for realizing desired performance and scalability.

> Relational and NoSQL InMemory Databases enable the otherwise difficult confluence of Horizontal scalability, ACID compliance and high performance at scale, that too with great architectural simplicity.

This ACID tradeoff in different schemes has been normative to the NoSQL database domain thus far. However, with the application of IMC technologies, emerging DB vendors are reclaiming this tradeoff and enabling the realization of NoSQL database qualities without compromising on ACID guarantees. Aerospike is one such emerging DB vendor.

Aerospike is flash-optimized, in-memory–ACID-compliant NoSQL database that in its own claims, reports the following performance characteristics.

1. Speed: 99% < 1 millisecond, predictable low latency,
2. Scale: 1 Million TPS on just a handful of servers
3. Reliability: 100 percent uptime with strong consistency (ACID).

While these claims have several subjectivities: around type of data, schematic and query complexities, access patterns, and hardware configuration, among others, Aerospike is today widely acknowledged as a leading, extremely high performance in-memory, ACID compliant NoSQL Database—a sharp realization of the possibilities that IMC technologies can enable.

3. Application servers and messaging middleware

Many enterprise applications are tier-based and consist of physical tiers such as data tier (relational database), business logic tier (application servers), web tier (web server), and integration/ messaging tier (messaging middleware). Now, to achieve high availability and higher processing capacity, all the tiers use clustered configuration.

Tier-based architectures have many shortcomings, for example:

a. They are hard to manage owing to multiple clustering models

b. They are difficult to install in virtualized environment because of the binding with static resources

c. They increase latency when transaction spans multiple physical tiers owing to network hops and disk writes (for reliability)

Now, to improve overall performance a common architecture practice is to add another tier: a caching tier. Caching addresses the performance and scalability problem for infrequently updated datasets but adds to the complexity. Further, caching solves limited use cases pertaining to read-mostly scenarios. It still does not solve the latency and scalability issues of write-mostly scenarios. That is where in-memory transactional application servers such as XAP (Gigaspaces) and ActiveSpaces (TIBCO) enable realization of extremely high-performance, distributed transactional applications but with limited complexity.

These in-memory transactional application servers allow placing data, logic, events, and transactions in memory and thus realize extreme low-latency architectures that scale well horizontally. These are known as space-based architecture (SBA). This contrasts with traditional N-tier architectures that add latency at each tier in addition to disk-drive input/output. This is another extreme possibility enabled with IMC.

4. Online analytical processing (OLAP) tools

OLAP cube is the core of the OLAP system. In response to complex business queries, OLAP cubes facilitate multidimensional data analysis and provide users with comprehensive

> **In-memory transactional application servers allow placing data, logic, events, and transactions in memory and thus realize extreme low-latency architectures that scale well horizontally.**

and valuable business insights in a short time (in a few seconds). The main issue with OLAP cubes created on disk-based databases is that they cannot deliver real-time analysis and reporting—something high performing businesses now expect—because OLAP cubes need to be updated in batches. For organizations with large data volumes, it results in unacceptable data latency, preventing business users from accessing current data and up-to-the-minute data analysis.

In-memory analytics and OLAP are useful to overcome such limitations. In-memory analytics eliminates the need to store pre-calculated data in the form of aggregate tables. It uses main memory for data storage which helps perform high-speed analysis compared to standard disk-based database systems, as these systems do not have to perform disk I/O to update or query data. Further, because the need for ETL and data warehouse is eliminated, reporting and analytics can be performed in real-time.

SAP HANA, for example, stores entire analytical database in memory and performs complex calculations over large data sets with blazing performance. Since OLAP cubes are not computed, in-memory OLAP tools, such as SAP HANA and TIBCO SpotFire, help deliver real-time BI.

5. Bigdata platforms

In last couple of years, data has grown exponentially. Many enterprises facing Bigdata challenges have readily embraced Hadoop, which has allowed them to store and process Bigdata in order of several hundreds of terabytes or even tens of petabytes.

> **BigStream processing is the next frontier in BigData. IMC platforms like GridGain are rapidly evolving to support these use cases and bring the potential of high performance In-memory computing to BigData stacks**

Hadoop provides distributed storage (HDFS) and massively parallel disk-based batch processing (Map-Reduce) infrastructure. Hadoop is suitable for performing offline batch analytics on Bigdata. However, the business demands are growing, it's no more only about being able to store and analyze extremely large data sets, but is increasingly also about being able to analyze streaming BigData sets as they arrive into the data system in conjunction with other existing large data sets, in near real-time.

A complementing set of Hadoop frameworks and distributed stream processing frameworks lead the stream analytics solution space today, however, IMC-based in-memory data grids (IMDGs) are an emerging and highly viable alternative that outpace most others.

IMDGs such as Gridgain use main memory as the primary data store. This 'memory-first-disk-second' architecture enables IMDGs to gain enormous speed that is much needed for real-time processing of BigData. Compared to disk storage, the available main memory storage here is very little—just a few hundreds of gigabytes. Therefore, IMDGs are designed to linearly scale on demand to hundreds of nodes. The data partition approach allows IMDGs to store several terabytes of data completely in memory. The distributed processing optimized for parallelism and data locality in IMDGs enables real-time processing and complex analytics on BigData.

Some of the leading vendors who offer IMDG: Red Hat (JBoss Infinispan), VMware (Gemfire), and Hazelcast, TIBCO (ActiveSpaces), and GridGain.

FACTORS AFFECTING IMC ADOPTION

Let us now assess the facilitating and inhibiting factors dominating the IMC play today.

Facilitating factors

1. **High velocity bigdata ingestion**

   Traditional disk-based systems are not capable of ingesting high-velocity BigData because of limitations with mechanical disks. Though some disk-based NoSQL database vendors claim that they can scale-out to ingest millions of messages per second, the fact is that they have to use very large clusters to accomplish the same. For example, in one of the benchmarks published by Aerospike (an in-memory NoSQL database vendor) showed that the insert throughput on 4-node cluster for Aerospike was approximately 275000 TPS whereas that for Cassandra and MongoDB (both disk-based NoSQL databases) was approximately 75000 and 25000 TPS, respectively.

   The massive shift to embrace problems at scale makes IMC solutions a great fit for an increasingly growing number of enterprise use cases and has been fuelling the demand in this space.

2. **Real-time analytics**

   Today, businesses are demanding real-time analytics in order to adapt to the changing business conditions and also to take advantage of short-window transient opportunities. Disk-based systems are incapable of simultaneously handling transactional and analytical workloads. Therefore, in the past, they were separated into different systems—OLTP for handling transactional workload and OLAP for handling analytical workload. The problem with this approach is that data is required to be transferred from OLTP systems to OLAP systems, which not only adds information access delays but also makes it less actionable.

   Here, in-memory computing platform gains ground as they are capable of simultaneously handling transactional and analytical workloads. Having single data management platform for OLTP and OLAP reduces overhead of multiple platforms, provides access to the latest updated data for analysis, and eliminates information delays. It gives an opportunity to perform analytics directly on OLTP data, thus allowing businesses to enable real-time analytics.

3. **Predictable high performance for unpredictable intensive workloads**

   High performance is indispensable for businesses operating on the Internet; even the slightest latency (in milliseconds) leads to loss of customer or revenue. For example, a leading e-commerce research claims that an additional latency of 100ms can reduce sale up to 1 percent or even more. Moreover, businesses running on the Internet such as e-commerce sites are required to interface with millions of users. Any business promotional campaign can draw huge traffic resulting in intensive and unpredictable workload.

IMC-based solutions form a great fit to such performance sensitive use cases. The memory-based system yield high performance, linear scalability, and ensure predictable performance (even under intensive workload) and dynamic scalability. This ensures elasticity (adding and removing capacity as needed). True linear scalability is not possible with disk-based system because generally accepted speed improvements with additional processors get restricted by overheads (such as IO waits in disk-based system) that inhibit parallelism. -Given limited disk orientation, IMC solutions stand out on this count as well and are expected to see extensive adoption in such high performance segments.

> InMemory computing based systems yield high performance, linear scalability, and ensure predictable performance and dynamic scalability.

with unpredictable workloads. True linear scalability is not possible with disk-based system because as proven using Amdahl's argument that speed improvement with additional processors gets restricted by any overheads (such as IO waits in disk-based system) that inhibit parallelism.

Inhibiting factors

1. Lack of standards, skills, and best practices

The main hindrance to widespread adoption of IMC is absence of any standards in this space. Currently, IMC vendors are developing their product without adhering to any standards. The lack of standard APIs and protocols is posing difficulties when it comes to application portability and interoperability. For example, if you have an expert team working on Gigaspaces IMDG, it is not easy to turn them into an expert in Gridgain or Hazelcast. Therefore, experts strongly believe that these issues are slowing down adoption of IMC and adoption would have moved faster for IMC if commonly agreed standards were available.

2. Fewer open-source vendors

While commercially backed solutions are a key enabler for widespread adoption of technologies in enterprises, open-source distributions have proven their worth as a massive force in fostering the early adopters ecosystem for ground breaking and assumption shattering technologies. They allow early adopters and business with unique needs and thus higher risk appetite to build early success stories and endorsements that can massively simplify the adoption path and perceived risk with emerging technologies.

Given IMC technologies are a significant offshoot of traditional solution approaches, especially in the technology world where much of our technology and architectures have grown around the constraints (both real and historical) in hardware and computer architectures, IMC adoption in the technology community has been hindered by conspicuous absence of open-source distributions of IMC solutions in the past.

In the past, IMC had to reckon with this issue. But, the future seems to present a different picture. Leading IMC vendors including GridGain and Aerospike have recently open sourced their platforms and are encouraging businesses to experience the value and power of these solutions without massive purchase processes and commercial risks.

3. **Limited Skill availability**

IMC is a relatively new paradigm in active development. Availability of trained engineering staff and seasoned consulting professionals are thus a challenge at this time. Given the many significant unlearning that these stacks demand, lack of general standards and absence of proven architectural patterns, this stack suffers the typical challenges with radical technologies during early adoption phases.

With the emergence of successful vendors, increasing network of implementation partners and the recent open-source momentum, the skill availability challenges are expected to drop rapidly, however for the interim IMC continues to be a niche technology and skill availability continues to be tight.

Conclusion

Leaving apart the teething challenges that are rapidly getting out of way, Cybage assesses that IMC technologies are mature and near-ready for primetime. They are already being widely adopted in niche segments and niche use cases and are now also starting to see rapid growth in general application segments where they bring the promise of not just performance and scalability but also massive simplification of architectures and extensive reduction in the time for implementation.

The parallel evolution in memory channel storage and spiraling cost optimization in primary memory manufacturing are driving IMC solutions to greater cost and scale accessibility.
Cybage looks at IMC technologies as a great opportunity, especially for enterprises whose success is linked to scale and performance.

No.	Acronyms	Full Form
1	IMC	**I**n-**M**emory **C**omputing
2	IT	**I**nformation **T**echnology
3	HANA	**H**igh-Performance **AN**alytic **A**ppliance
4	RDBMS	**R**elational **D**ata**b**ase **M**anagement **S**ystem
5	XAP	e**X**treme **A**pplication **P**latform
6	SQL	**S**tructured **Q**uery **L**anguage
7.	OLTP	**O**n**L**ine **T**ransaction **P**rocessing
8.	IMDG	**I**n **M**emory **D**ata **G**rid
9.	OEM	**O**riginal **E**quipment **M**anufacturer
10.	SSD	**S**olid **S**tate **D**rive
11.	IMDB	**I**n **M**emory **D**ata**b**ase
12.	DDR	**D**ouble **D**ata **R**ate type 3
13.	DIMM	**D**ual **I**nline **M**emory **M**odule
14.	HTAP	**H**ybrid **T**ransactional & **A**nalytical Processing **P**latform
15.	SCP	**S**upply **C**hain **P**latform
16.	ACID	**A**tomicity, **C**onsistency, **I**solation, **D**urability
17.	NOSQL	**N**ot **O**nly **SQL**
18.	BI	**B**usiness **I**ntelligence
19.	OLAP	**O**n**l**ine **A**nalytical **P**rocessing
20.	GPU	**G**raphics **P**rocessing **U**nit
21.	CEP	**C**omplex **E**vent **P**rocessing
22.	DRAM	**D**ynamic **R**andom **A**ccess **M**emory
23.	WAL	**W**rite **A**head **L**ogging
24.	IMRDB	**I**n **M**emory **R**elational **D**ata**B**ase
25.	DB	**D**ata**B**ase
26.	SBA	**S**pace **B**ased **A**rchitecture
27.	ETL	**E**xtract **T**ransformation **L**oad
28.	HDFS	**H**adoop **D**istributed **F**ile **S**ystem
29.	TPS	**T**ransactions **P**er **S**econds
30.	API	**A**pplication **P**rogramming **I**nterface
31.	IO	**I**nput **O**utput

Cybage Software Pvt. Ltd.

[An SEI-CMMI Level 5 & ISO 27001 Company]

HQ: Cybage Towers, Survey No 13A/ 1+2+3/1, Vadgaon Sheri, Pune 411014 |
Tel: 91 20 6604 4700 | Fax: 91 20 6604 1701
Pune | Hyderabad | Gandhinagar | Seattle | New Jersey | San Francisco |
Atlanta | Austin | London | Frankfurt | Amsterdam | Sydney
www.cybage.com

CYBAGE
Delivering Value. Scientifically.