



**INTESA
SANPAOLO**

CASE STUDY

500 Years of Banking Innovation Continues with Software-Defined Memory

Banca Intesa Sanpaolo: 500 Years of Banking Innovation

The origins of Banca Intesa Sanpaolo go back to 1563 in Turin, when a brotherhood known as “Compagnia di San Paolo” laid the foundations for what a few years later would become the Monte di Pietà (Mount of Piety), an institutional pawnbroker run as a charity in Europe from Renaissance times until today. This new institution innovated by offering loans at extremely low interest rates to provide the poor with an alternative to usury.

At the end of 1991, credit activities were delegated to a newly established joint stock company, the San Paolo Bank of Turin SpA. At the same time, the business related to public interest and social utility remained in the “Old San Paolo” and took on the name of Compagnia di San Paolo in homage to its original history.

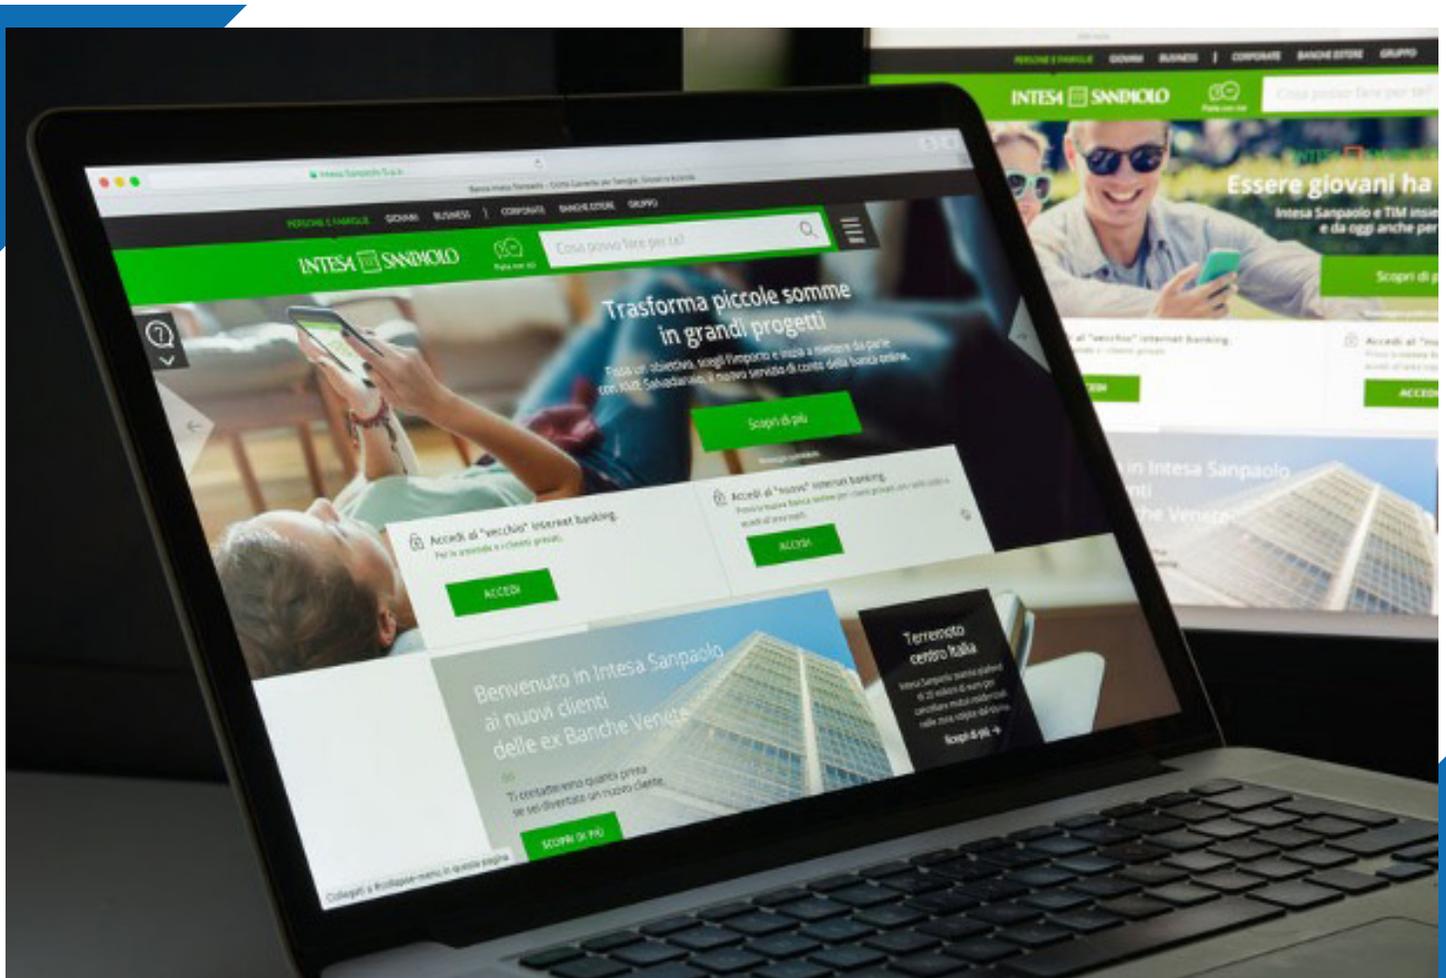
Today, Banca Intesa Sanpaolo is the leading banking group in Italy and one of the top banking groups in Europe, with a market capitalization over 37 billion Euro as of November 2020. Intesa Sanpaolo has over 5,000 branches and 40 million customers in approximately 40 countries worldwide.



Online Banking Drives Migration to Container-Based Microservices

To continue their heritage of innovation that started almost 500 years ago, Banca Intesa Sanpaolo launched an initiative to undergo a digital transformation. The goal was to sharpen the bank's competitive edge in online banking by reducing financial application development time and footprints for greater agility, while improving scalability and availability.

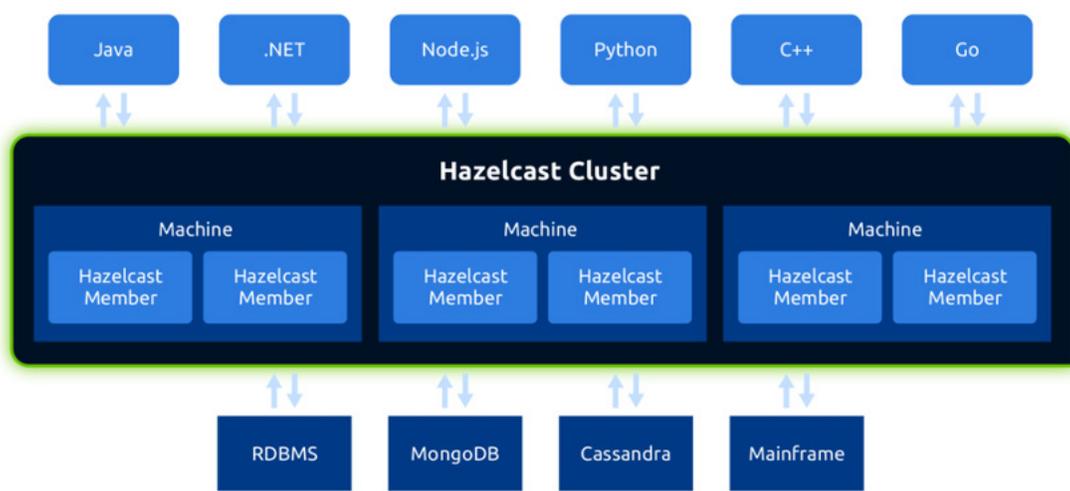
To achieve their goal the bank embraced a container-based microservices architecture and a migration from monolithic to multi-tier applications. It was also critically important for the IT organization at Intesa Sanpaolo to integrate any modern technologies without impacting the stability of their core systems.



Hazelcast Powers Home Banking App Scalability

Hazelcast in-memory data grid was deployed to power their new containerized home banking microservices. By allowing access to a shared pool of memory across a cluster of servers, Hazelcast allows the new generation of apps to access data in-memory lightning fast, as well as scale quickly and smoothly as the in-memory data sets grow.

Hazelcast Architecture



An IMDG (in-memory data grid) is a set of networked/clustered computers that pool together their memory to let applications share data structures with other applications running in the cluster. The primary advantage is speed, which has become critical in an environment with billions of mobile, IoT devices and other sources continuously streaming data. With all relevant information in memory in an IMDG, there is no need to traverse a network to remote storage for transaction processing. The difference in speed is significant – minutes vs. sub-millisecond response times for complex transactions done millions of times per second. Source: Hazelcast

Banca Intesa Sanpaolo data center in Turin



Persistent Memory Cuts Cost in Half and Doubles Capacity/DIMM

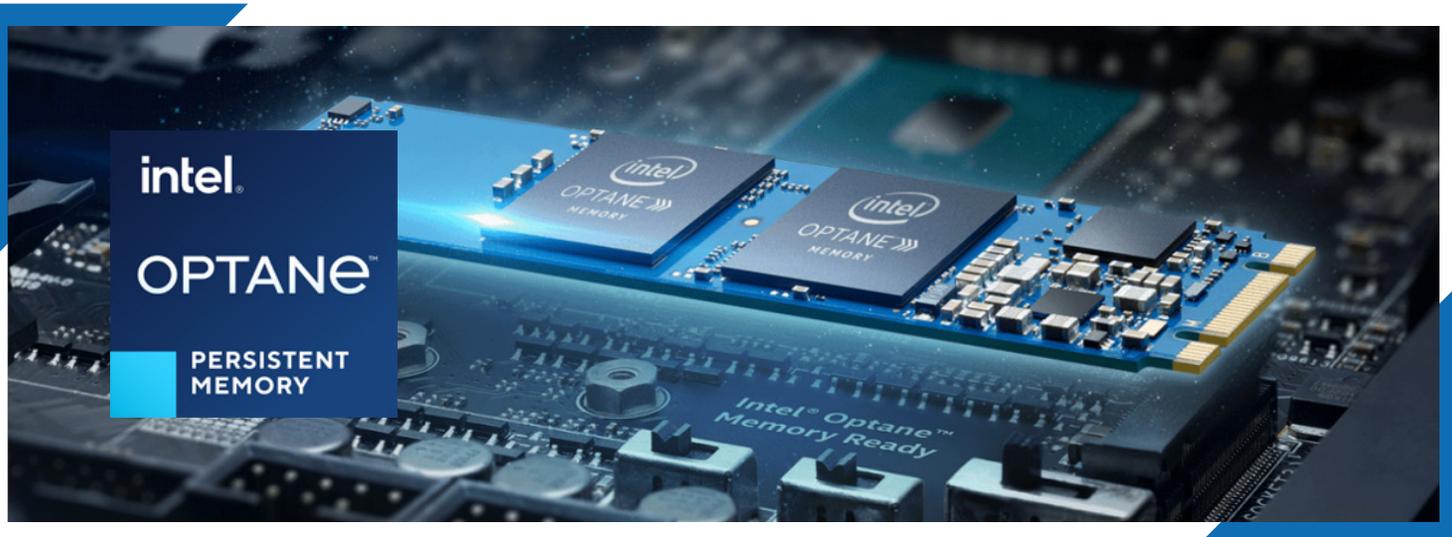
For in-memory computing environments, memory can be 70% of the cost of the application servers. As the size of the data increases, the need to increase memory capacity and lower the cost of expensive DRAM becomes a priority. Banca Intesa Sanpaulo addressed the need for greater capacity, and the growing cost of memory, by adding Intel® Optane™ Persistent Memory (PMEM).

With a capacity of 512GB per DIMM, PMEM offers twice the capacity of the largest 256GB DRAM DIMM. At the same time, the cost of a PMEM DIMM is as much as one-third less than a comparable DRAM DIMM. With PMEM, dual-socket servers with 6 DIMM slots per socket can support up to 6TB of memory.

The greater the memory footprint, the longer it takes to recover data from storage after a server crashes, power is out, or after any other event that causes a server to go down.

| Capacity | PMEM | DRAM |
|-----------|------------|------------|
| 1 x 512GB | \$13.86/GB | - |
| 1 x 256GB | \$7.02/GB | \$18.94/GB |
| 1 x 128GB | \$4.00/GB | \$13.67/GB |
| 1 x 64GB | - | \$7.65/GB |
| 1 x 32GB | - | \$8.43/GB |
| 1 x 16GB | - | \$9.37/GB |

The IT team at Intesa San Paulo wanted to address the growing memory blast zone by leveraging the persistence of PMEM and other memory services for high availability. This led them to MemVerge and the new world of Big Memory Computing consisting of Big Memory HW + Big Memory SW.



Big Memory Computing

The next generation of in-memory computing, Big Memory Computing transforms low-density, expensive, and volatile DRAM-only environments, into high-density, lower-cost, and highly available environments based on DRAM + persistent memory (PMEM).

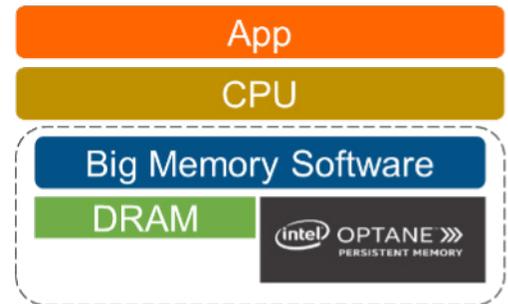
Big Memory Hardware is any type of server memory. Today that is DRAM plus higher density and lower cost Intel® Optane™ Persistent Memory (PMEM).

Big Memory Software virtualizes DRAM and PMEM, creating a pool of software-defined memory and delivering software-defined memory services for higher performance, availability, and agility.

Big Memory Software also serves as an abstraction layer that allows all applications in a data center to benefit from new types of memory, memory interconnects, processors, memory allocators, etc., all without changes to the apps.

Memory Machine™ from MemVerge is the world's first Big Memory Software and Banca Intesa Sanpaulo deployed the software to transform their in-memory grid into a high-availability data storage tier.

Big Memory Computing



Once DRAM and PMEM are virtualized, they become a pool of software-defined memory. Big Memory Software then serves as a platform for memory services that tier data for higher performance and provide memory-to-memory snapshots for instant recovery and high-availability.



Opening the Door to Big Memory



A Foundation for Software-Defined Memory Infrastructure

Once Memory Machine was installed and the DRAM and PMEM was virtualized, the software provided Banca Intesa San Paolo with a foundation for software defined memory infrastructure that required no changes to their apps, and in the years ahead will deliver:

1. **Higher PMEM Performance** that opens the door to lower cost persistent memory. Without Memory Machine, Intel Optane Persistent Memory runs slower than DRAM. With Memory Machine, the pool of DRAM and PMEM run as fast or faster than DRAM. This enables mass deployment of lower cost and higher density PMEM without sacrificing performance.
2. **Higher Availability.** Memory Machine for the 1st time transforms memory into an HA environment. Most memory is not protected with snapshots and replication. The large memory footprints that are protected take minutes to snapshot and hours to recover. Memory Machine enables instant snapshot and recovery, decreasing business disruption by 1,500x.
3. **On-Prem & Public Cloud.** For hybrid cloud environments, Memory Machine software supports container-based microservices both on-prem and in the public cloud.



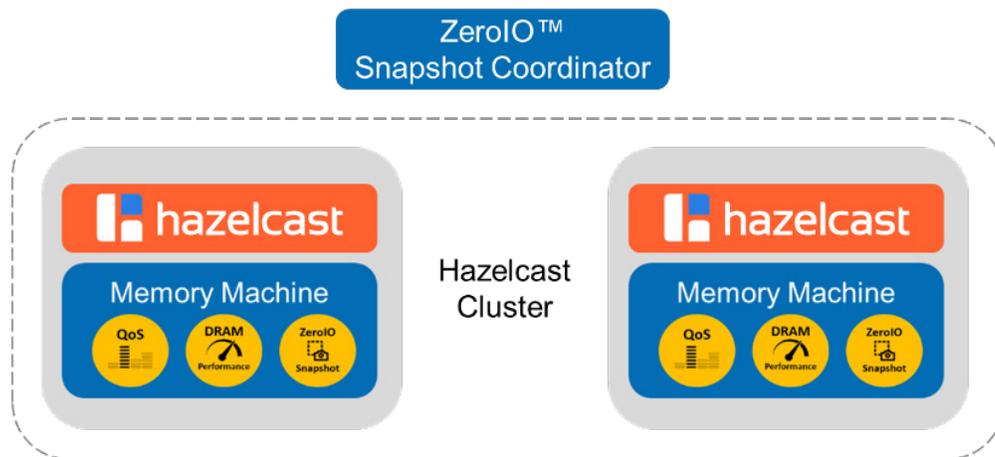
Memory Machine™
The World's First Big Memory Software

Memory-to-Memory Snapshot & Recovery in a Hazelcast IMDG

Memory Machine features the world's first memory-to-memory snapshot technology. This ability to snapshot data from DRAM to persistent memory transforms the volatile low-availability memory tier into a higher availability tier.

But Banca Intesa San Paolo needed more to be at the forefront of responsive, 24x7 online banking. They needed the ability to send the memory-to-memory snapshots to other server nodes in a Hazelcast cluster. Then in the event one member of the cluster was down, the other members could quickly recover the data.

MemVerge responded by providing a Snapshot Coordinator that allows snapshots to be transported to other servers so they can restore a database instance from PMEM. Recovering a database instance from storage that used to take minutes to hours, now takes approximately 2 seconds with software-defined memory-to-memory services.



The 500 Year Tradition of Innovation Lives On

Since DRAM was invented in 1969, the server memory model has changed little as DRAM continues to be expensive, volatile, and with higher capacities achieved only by constant IO to slower storage used as an extension of memory.

The 500-year tradition of innovation lives on as Banca Intesa Sanpaolo moves to the forefront of an IT megatrend towards Big Memory Computing that will cost-effectively allow all data to live in-memory. By deploying Big Memory, the bank is uniquely positioned to process massive data sets at real-time speeds and avoid downtime with memory-to-memory data protection services.

Looking forward, Banca Intesa Sanpaolo is looking to Big Memory and Memory Machine software for real-time machine learning applications. With tens of millions of transactional and behavioral events per day, their machine learning architecture needs to be fast and resilient for their real-time recommendation engines to promote offers to customers when specific conditions are met.

Learn More

[Everything you need to know about Big Memory in 3 minutes](#)

[MemVerge Corporate Brochure](#)

[The Skinny on Memory Machine](#)

[IDC Big Memory Definition and PMEM Forecast Presentation](#)

[Big Memory AI/ML Solution Brief](#)

[Tech Field Day, Big Memory Architecture and Use Cases](#)

[Demo: Creating Clones of Redis VMs in Microsoft Azure](#)

[Demo: Memory Snapshots and Managing from GUI and Command Line](#)

[Demo: Cloning an 800GB kdb+ Database in Seconds](#)

[Demo: See kdb+ in-memory on AWS run faster with Memory Machine](#)



[Big Memory Impact on Banking](#)

Nicola Carotti

Head of Cloud and Hybrid Datacenter
Banca Intesa Sanpaolo



What happens in memory,
stays in memory...