

vSMP MemoryONE for EC2

Large-memory instances
at a fraction of the cost



DATA SHEET

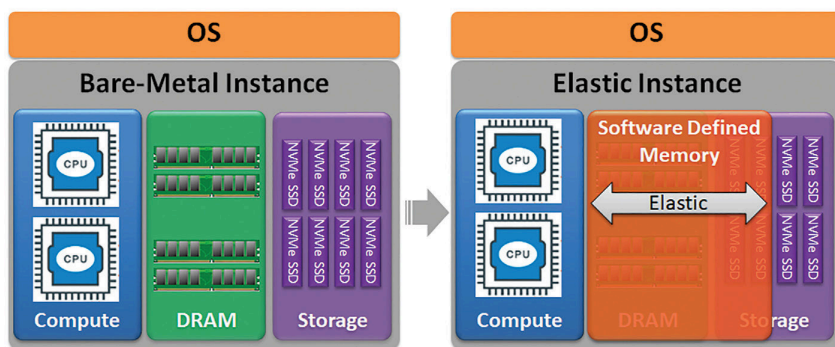
MEMORY EXPANSION FOR THE CLOUD

Organizations are leveraging in-memory computing to drive superior application performance and obtain meaningful insights using advanced business analytics.

In-memory computing can be limited by a set of constraints on the amount of memory available per server or cloud instance. This is due to modern server architecture, as well as prohibitive DRAM pricing.

vSMP MemoryONE for EC2 expands system memory of NVMe-based instances by leveraging NVMe SSDs to transparently be used as memory.

AWS users benefit from much larger system memory and more economically efficient instances with DRAM-like performance.



HIGHLIGHTS

- Largest **on-demand** memory instances
- Up to **70%** savings
- Billed **hourly**, with no yearly commitment required
- **Transparent** to OS and applications
- **No OS changes**, supports standard linux AMIs
- **No application changes** required

USE CASES

- In-memory databases
- Business analytics
- Data warehousing
- Application caching
- Container-based environments
- Genomics & AI research
- Electronic design automation
- IoT time series analysis

ELASTICALLY SCALING MEMORY AND MEMORY/CPU RATIO

In-memory databases or high-concurrency environments, such as container-based applications, are examples of cases in which memory usage can quickly outpace processing capabilities. At this point, options are limited and typically require the addition of more instances (scale out) to provide the extra memory, or the use of special larger systems, which are not typically available as an ad-hoc/on-demand resource.


vSMP MemoryONE takes advantage of cost-effective EC2 instances (available on demand or on Spot), while expanding the available system memory of those instances to multiple TBs, as well as increasing the memory/CPU ratio.

In addition, MemoryONE is using the latest-generation instances, bringing additional CPU power.

COST SAVINGS

EC2 instances combined with vSMP MemoryONE for EC2 can provide more memory, while also reducing cost. For example, an AWS i3en.metal with vSMP Memory ONE provides up to 6TB of system memory, saving up to thousands of dollars per month per instance for on-demand instances. Savings can be even higher if using Spot Instances.

System Memory vs. Hourly Cost



vSMP MemoryONE users benefit from very large memory instances on-demand, with a significantly lower hourly cost, and without any long-term commitment to proprietary hardware or service.

ORDERING AND AVAILABILITY

vSMP MemoryONE for EC2 is available as an EC2 AMI, on [AWS Marketplace](#).

SAMPLE APPLICATIONS

- In-Memory databases
- Apache Spark
- Redis open-source
- memcached
- KVM
- Docker
- MongoDB
- Cassandra
- MySQL
- MariaDB
- SOAPdenovo
- Abaqus Standard
- Ansys Mechanical
- Altair Optistruct

PERFORMANCE

The performance chart shows the results of a segmented-SGEMM benchmark (single precision floating general matrix multiply) – a form of computation that is used in artificial intelligence applications. We compared the compute efficiency of a r5.metal DRAM-only instance (with 768GB of memory) running a 689GB workload to the **same instance** with vSMP MemoryONE expanding memory by 4X (a total of 3,072GB of memory), running a 2,775GB workload.

As the chart clearly shows, the **computational efficiency is at 93%** compared to a 4X smaller DRAM-only benchmark.

SGEMM:

Performance vs. Workload memory use

