

CASE STUDY

Memory Machine™ Cloud

WaveRider feature continuously right sizes BLAST workload to slash compute costs and accelerate completion time.

Memory Machine WaveRider Reduces BLAST Compute Cost by 91% and Accelerates Run Time by 31%



Hongsonbio focuses on applications in the fields of biomedical technology, life science research and human health, and is committed to providing one-stop, all-round services and system solutions for clinical medicine and life health. Its laboratories have been established in Tianjin and Shanghai, with a total laboratory area of over 5,000 square meters.

BLAST Workload

One of their primary computational analytics tools is BLAST (Basic Local Alignment Search Tool), which is used to find regions of similarity across nucleotide and protein sequences. Honsun runs most of their BLAST searches on cloud computing resources.

Spikey Use of CPU and Memory Resources During Runtime

BLAST workloads can be “spikey” with the need for cloud resources such as virtual CPUs (vCPUs) and memory varying significantly throughout a run, as shown in the graph below.


No Single Compute Instance is the Right Size

The simplest solution to this is to provision application for the largest resource spikes (essentially overprovisioning).

A representative BLAST run

overprovisioned for 32vCPUs and 64GB of cloud memory completes in 11 hours and 44 minutes, with a an on-demand cost of \$17.22. Under-provisioning (provisioning for average resource usage) is an alternative and reduced the on-demand costs to \$12.63, or 27% less expensive. But under-provisioning increased run time to 18 hours and 4 minutes, a 53% increase in runtime. Running BLAST in a Spot instance can significantly reduce costs, but a spot evacuation would stop the job, requiring it to be re-run from the start, impacting overall productivity.

Spikey CPU Usage by a BLAST Workload



Right Sizing During Runtime with Memory Machine WaveRider

Memory Machine Cloud includes a feature called WaveRider which is the answer to the “spikey” cloud workload issue. Patented AppCapsule technology allows long-running BLAST applications to run safely on cloud spot instances because they can surf to a new compute instance when the Spot is reclaimed. WaveRider extends this ability by allowing applications to right size to more or less resources as needed during runtime.

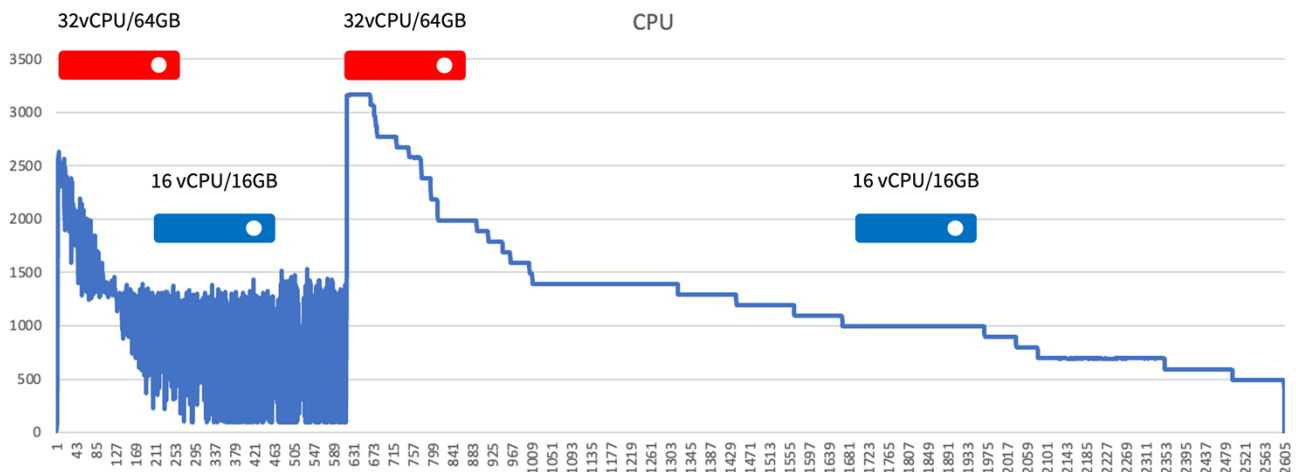


How WaveRider Worked at Honsun

The graph below shows how WaveRider enabled a HonSun BLAST workload to continuously and automatically right size compute resources. WaveRider found larger and smaller cloud instances and moved the application between these instances as the needs of the workload changed.

The migration between instances is fully automated and does not require user intervention. The automation is triggered through user-created rules that are built using Memory Machine’s Float GUI and command language. This enables WaveRider migrations to be tuned as necessary to optimize application cost and performance.

Spikey CPU Usage by a BLAST Workload



The Results: Significant Savings & Reduced Run Times

The chart below provides a comparison between the different scenarios: OnDemand with 32vCPUs/64GB RAM; OnDemand with 16vCPUs/16GB RAM; spot with 32vCPUs/64GB RAM protected by Memory Machine; spot with 16vCPUs/32GB RAM protected by Memory Machine; and auto VM sizing using Memory Machine WaveRider.

WaveRider provides the perfect balance between performance and cost, with a cost reduction of 93% when compared to the large, 32vCPU, 32GB RAM, OnDemand option, but with nearly the same performance of 12h 23m execution time vs 11 h 44m. For organizations running 1,000 BLAST applications per day in the cloud, using Memory Machine with WaveRider would result in a savings of \$16,065 per day, or \$4,173,000 per year.

Examples of Savings & Faster Completion Time with WaveRider

Use Cases	Time	Delta	Cost	Delta
OnDemand (32 vCPU, 64GB)	11h 44m	Fastest	\$17.22	Highest cost
OnDemand (16 vCPU, 16GB)	18h 4m	53% slower	\$12.63	27% less cost than 32vCPU
Memory Machine (Spot; 32 vCPU, 64GB)	11h 44m	Fastest	\$1.93	89% less cost than 32vCPU
Memory Machine (Spot; 16 vCPU, 32GB)	18h 4m	53% slower	\$1.48	91% less cost than 32vCPU
Memory Machine with WaveRider	12h 23m	31% faster than 16vCPU	\$1.17	93% less cost than On-demand

Memory Machine Cloud

Memory Machine Cloud Edition from MemVerge is a Cloud Automation 2.0 Platform. The software visualizes compute and memory resources and usage, so users can easily see opportunities to right-size and save. Memory Machine then automates complex deployment of a single, to hundreds of job runs. And the software decouples workloads from VMs so they can surf to larger and smaller resources based on the need of the workload; provides observability into resource usage.

Available in the AWS Marketplace

Memory Machine Cloud Edition today is certified as Spot Ready and available in the AWS Marketplace.

Get Started Free

[Request a demo](#) and a MemVerge expert will explain how Memory Machine works, then walk you through a demonstration that shows how easy it is to deploy and how powerful the results are.

[Request a free trial](#) and a MemVerge expert will set you up, or we can provide access for you to play in a Sandbox set-up in AWS.

Resources

[Memory Machine Cloud Edition Landing Page](#)

[Memory Machine Cloud Edition Data Sheet](#)

[Memory Machine Cloud Edition QuickStart Guide](#)

[Memory Machine Cloud Edition User Guide](#)

[Memory Machine Cloud Edition QuickStart Video](#)

[How to use Memory Machine Cloud Edition Video](#)

