MemVerge

Accelerate Scientific Research and Easily Optimize Cloud Costs

Challenges

Running scientific workflows and performing data analysis in the cloud can be time consuming, unpredictable, and expensive

For R&D organizations using the cloud, it is often a tale of two cities. On the one hand, the cloud promises agility via compute resources at your fingertips, scalability via seemingly unlimited resources, and lower cost in that you only pay for what you use. The reality is many barriers to entry still exist to realize these promises. In addition to the learning curve that comes with managing cloud infrastructure, it is difficult to predict the resource requirements of various research computing workload and even more challenging to take advantage of attractive cloud pricing models such as EC2 Spot at scale.

The MemVerge Solution

Easily turn AWS into a next-generation super computing environment in minutes, accelerate scientific research, and optimize cloud costs

MemVerge's Memory Machine Cloud "MMCloud" is a powerful and intuitive container orchestration platform for running data-intensive pipelines and interactive computing applications on AWS. The platform abstracts away DevOps complexity and makes it easy to safely run workloads on EC2 Spot, automatically right-size EC2 resources, and save cloud costs. MMCloud seamlessly integrates with popular workflow managers like Nextflow and runs in your own AWS account, ensuring data never leaves your IT approved environment, and you are in complete control of your cloud costs.

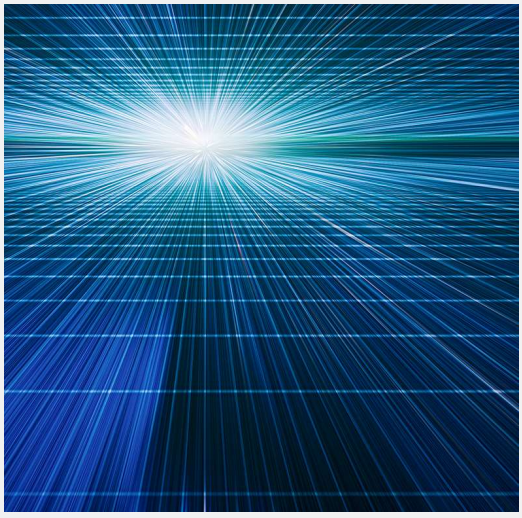

Benefits

MMCloud benefits the customer in a multitude of ways, through automating job execution in the cloud, through accelerating scientific research, and through optimizing cloud costs.

Automate Job Execution

MMCloud reduces the need for a typical DevOps team of 3-5 FTEs to scale, orchestrate, maintain and support numerous scientific workflows. In addition to automated cloud resource management, provisioning, and deprovisioning, detailed cloud resource utilization metrics are available in real-time and as reports to help research teams optimize their workflows.

Accelerate Scientific Research

MMCloud has shown the ability to improve wall times by 30-40% while at the same time optimizing cloud compute costs by right-sizing cloud resources at runtime. By automating checkpoint and recovery on EC2 Spot, jobs never need to start over from the beginning.

Optimize Cloud Costs

MMCloud puts you in the drivers seat when it comes to how to optimize cloud costs. Researchers can optimize for maximum savings by using a "Spot only" VM policy, thus saving upwards of 70-80% on compute costs vs. using EC2 On-Demand. Researchers can also optimize for both performance and cost by using a "Spot first" VM policy combined with an auto right-sizing policy to ensure resources are never over or under provisioned.

(MemVerge Contact: Jing Xie (jing.xie@memverge.com)

MemVerge on AWS

The benefit of MemVerge's partnership with AWS is two-fold. First, MMCloud enables customers to leverage the compute & storage resources in their own AWS account, enhancing security & data privacy, and improving connectivity to other AWS resources. In addition, customers benefit from economies of scale by consolidating all genomics analysis through their own accounts. Second, the MMCloud software is available through the AWS Marketplace where customers can easily spend their commits in a smart way and implement quickly. MMCloud provides both an easy-to-use GUI and a CLI based interface to enable a simple and effective user experience for all researchers.

MemVerge

Case Study: TGen

Challenges

- High failure rates running large Nextflow workflows on EC2 Spot made it difficult to reliably and cost effectively run on the cloud
- No easy way to right-size and optimize cloud resources for a large variety of computational pipelines that needed to be run
- Low confidence to run on the cloud due to poor visibility of performance and cloud resource utilization of existing workflows

Solution

Bioinformaticians at TGen used MMCloud as the computing environment for executing large Nextflow workflows on EC2 Spot.

Results

Bioinformaticians at TGen enhanced workflow execution with EC2 Spot and reduced failure rates from as high as 80% to near zero. MMCloud also optimized and reduced wall times further by automatically right-sizing EC2 instances at runtime, preventing out of memory errors that would normally result in reruns. Lastly, MMCloud provided organization-wide visibility into job-level resource utilization reports and insights.

Features

Job Automation with Built-In Checkpoint & Recovery

It is now safe, effective, and easy to run complex data pipelines on EC2 Spot instances, saving upwards of 70-80% on compute costs vs. EC2 On-Demand. If a Spot instance is reclaimed, MMCloud automatically checkpoints and migrates the running job to resume on a new EC2 instance versus having to restart from the beginning. Launch next-gen-sequencing workflows at night on EC2 Spot and sleep soundly, MMCloud is managing it for you.

Automated Right-Sizing Matches Your Workload to the Compute

Running a new data pipeline in the cloud just got easier and more cost efficient. MMCloud will supersize or downsize the EC2 instances at run-time based on real-time resource utilization metrics and auto-migration policies you define. Say goodbye to "OOM" out of memory errors and slow job execution with MMCloud's right-sizing service.

Visit AWS Marketplace to purchase or start a Free Trial today.

Get started with MemVerge solutions on AWS

aws Qualified Software

aws PARTNER Amazon EC2 Spot Ready

(MemVerge Contact: Jing Xie (jing.xie@memverge.com)