

wenovate Cloud Cost Optimization and Spot Fleet Comparison

Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity.

Spot Fleet

A Spot Fleet is a collection, or fleet, of Spot Instances. The Spot Fleet attempts to launch the number of Spot Instances that are required to meet the target capacity that you specified in the Spot Fleet request. The Spot Fleet also attempts to maintain its target capacity fleet if your Spot Instances are interrupted due to a change in Spot prices or available capacity.

wenovate Cloud Cost Optimization

wenovate Cloud Cost Optimization first uses predictive algorithms to predict Spot behavior, capacity trends, pricing, and interruption rate. Whenever there's a risk of interruption, wenovate Cloud Cost Optimization acts accordingly to seamlessly balance capacity, ensuring 100% availability and no risk of downtime. This means that your application will always run on the most cost-efficient collection of instances – the best-priced Spot Instances when available and falling back to on-demand when not. In addition, prioritizing any reserved instances you may already own.

Feature	Spot Fleet	wenovate Cloud Cost Optimization
SLA	No SLA for Spot Instances. The Instances can be terminated within 2 minutes notice	With wenovate Cloud Cost Optimization, your capacity is guaranteed and backed by an SLA.
Prediction	When Amazon EC2 interrupts your Spot Instance, it emits an interruption notice, followed by a termination of your instance. The maximum notification time you will receive from AWS is 2 minutes.	Based on historical and statistical data, wenovate Cloud Cost Optimization predicts interruptions ahead of time -- up to 15 minutes ahead -- and automatically relocates instances into another Spot capacity pool. In addition, wenovate Cloud Cost Optimization uses its

		<p>predictive technology every time it launches a new Instance, to make sure that instances will be placed into the most available & low-priced capacity pools.</p>
Instance Auto Recovery	Not Supported. Spot fleet can pause an instance and hibernate it until it can be available only.	<p>wenovate Cloud Cost Optimization will predict Spot instance termination and automatically replace the terminated Instances with the <u>best available Instance type across multiple AZs</u></p>
Automatic Fall Back to On-Demand	Not Supported. If the Spot market is unavailable the Instances will be terminated or hibernate until the market will be available again	<p>Once wenovate Cloud Cost Optimization identifies that all the Spot market are unavailable, the default behavior is to launch an On-Demand Instance instead</p>
Determination of instance types to launch	Determine the instance type according to the EC2 Fleet strategy	Determine instance types and pricing model intelligently, based on the workload performance requirements, Spot capacity-pool availability, and price
3rd party & AWS services integrations	<p>Automatic integration with the following services:</p> <ul style="list-style-type: none"> ● EMR ● ECS ● CloudFormation ● Auto Scaling Groups 	<p>Automatic integration with the following services:</p> <ul style="list-style-type: none"> ● Chef ● Nomad ● CodeDeploy ● OpsWorks ● Rancher ● Right Scale ● Mesosphere ● Kubernetes ● Docker Swarm

		<ul style="list-style-type: none"> • ECS • AWS Batch • Beanstalk • EMR • RDS • CloudFormation
Automatic Backup	No Automatic backups	Support for Stateful applications using scheduled snapshots of your AMI and attached EBS volumes. With the Auto Backup feature, you can maintain data persistence within your cluster. In the case of any instance replacement, wenovate Cloud Cost Optimization will use the last snapshot recorded according to the defined interval.
EBS Volumes	Partially Supported, only for Manual Instance Shut Down, but not for Spot termination.	Full Support - "Hot EBS migration". When a recovery occurs the system will automatically attach any external volume of the previous instance and it will continue service from the same state.
Containers Autoscaling	EC2 Spot Fleet & Auto Scaling group is using traditional scaling metrics to autoscale Containerized applications.	For Containerized applications, wenovate Cloud Cost Optimization is using its Autoscaler to determine the best blend of instance types from different instance families and sizes (including GPU & FPGA support) to match different containers requirements (CPU, Memory, Network, etc.) in the cluster.

Graceful draining	When Amazon EC2 interrupts your Spot Instance, it emits an interruption notice, followed by a termination of your instance	wenovate Cloud Cost Optimization predicts interruptions ahead of time and automatically relocates instances into another Spot capacity pool. wenovate Cloud Cost Optimization makes sure that the preemption is done gradually, maintain TCP/HTTP state and Containers constraints to ensure service uptime
RI Utilization	EC2 Auto Scaling group will launch exactly the amount of instances you configured for each pricing model, without caring for underutilized RIs	wenovate Cloud Cost Optimization makes sure to first utilize your Reserved Instances, before launching Spot or any other On-Demand instances
Maintaining Private IP	Not Supported	Supported
Vendor Support	AWS	AWS, Azure, GCP