

SEDAI FOR AWS LAMBDA

Autonomously Optimize Serverless Functions



Achieve the highest levels of performance, availability and cost efficiency for AWS Lambda functions with the Sedai Autonomous Cloud platform.

Results

75% Latency reduction

33% SRE workload reduction

Key Benefits

- **Improve Customer Experience through reduced latency**
- **Reduce cost through memory rightsizing and provisioned concurrency management**
- **Minimize SLO violations**
- **Improve SRE productivity by reducing toil**
- **Release with confidence using early insights into code performance in production**
- **Easy setup with agentless monitoring, integrations with popular monitoring and other tools**

The Challenge of Serverless Optimization

In dynamic production environments it has become virtually impossible for companies to continuously optimize serverless functions & overall application performance. The result is either application performance and latency challenges (and potentially lost revenue), or wasted cloud spend.

Even if a point-in-time optimization is made it can become outdated quickly due to constantly changing & unpredictable traffic loads and changes in the serverless function code base with every new release.

The task becomes unmanageable especially when teams run 100s or 1000s of serverless functions. SRE teams also must spread their bandwidth across multiple priorities.

To optimize serverless function teams must continuously:

1. Rightsize function memory
2. Manage provisioned concurrency
3. Choose how and when to warm up serverless functions and how to keep them warm
4. Trace errors in this heavily distributed environment

5. Manage optimization along with timeouts and reserved concurrency
6. Keep costs low

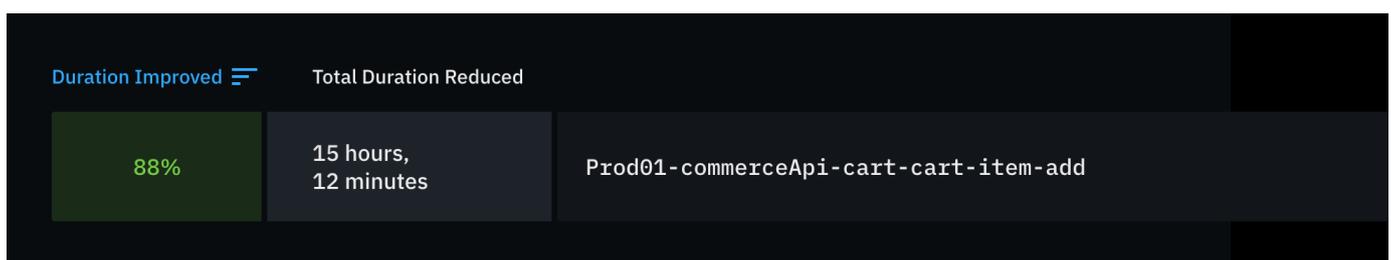
The Autonomous Approach

Sedai manages serverless functions autonomously using state-of-the-art machine learning and predictive engines that detect potential optimization opportunities and issues within your serverless functions and take autonomous action to address them in production, improving performance, ensuring availability, and optimizing cloud costs. Acting as an intelligent autopilot for SREs, Sedai eliminates significant toil for SREs so they can increase time spent on innovation. Sedai intelligently discovers AWS Lambda functions, analyzes traffic patterns and metric performance and then continuously manages your production environments — without manual thresholds or human intervention.

Autonomous Optimization

Improve Customer Experience through reduced latency

For most SRE teams, the number one serverless concern is latency. Latency drives customer experience and revenue. Sedai reduces latency by up to 95% at a function level by finding the optimal parameters or the function by optimizing memory settings (and indirectly CPU), concurrency (provisioned and reserved) and timeouts for Lambda functions. In the example below the execution time of a function that adds an item to an ecommerce shopping cart has reduced by 88%:



Sedai also reports on the number of days of cumulative latency reduction achieved and the cloud spend needed (where applicable) to achieve this.

Reduce Cloud Costs

Sedai improves your serverless costs through autonomous optimizations. When lambda functions are not extensively used (e.g., during night time traffic lows), Sedai scales down memory based on the traffic pattern. In addition to memory, Sedai also cost-optimizes provisioned concurrency to avoid over-provisioning and optimize function warm up. After finding improvements that benefit both cost and latency, Sedai gives you the option to then prioritize cost alone, latency, or make a balanced tradeoff.

Hit SLOs

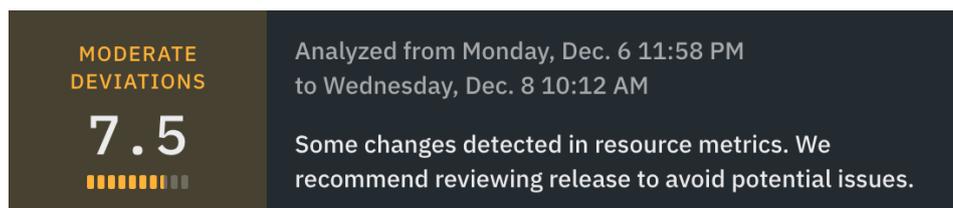
Sedai enables users to configure service level objectives (SLO). These SLOs are then autonomously managed by Sedai, which ensures that performance for each service falls within the SLO. For example, if a service slows down (below the latency SLO) during peak traffic, Sedai will scale it up to compensate and make sure it stays within the SLO.

Improve SRE productivity by reducing toil

Sedai enables you to support workload growth without needing to linearly scale up the SRE team. Sedai's autonomous detection and optimization means the work of detecting & resolving issues and manually finding optimization is shifted from SREs to Sedai which can take care of many tasks without their input. Sedai enables SREs to be less concerned with what is happening in the infrastructure and the chaos of alert fatigue and reactive troubleshooting, and instead focus on new deployments, site enhancements and other priorities. SREs now use Sedai to check new applications and services are onboarding effectively, that autonomous mode is on, and reviewing optimizations executed by Sedai.

Release with Confidence

Sedai provides early feedback on the performance of every serverless release with Release intelligence. Sedai will autonomously detect new Lambda releases entering production. Sedai analyzes the release compared to the prior release on latency and error count using models that take into account seasonality and traffic patterns under load. Each release is rated on a scale of 1-10 based on the amount of deviation from the prior release. You can receive recommendations on whether to rollback a release if it doesn't meet certain criteria or parameters. In addition, scores provide signals to developers on where effort should be placed to optimize the underlying application code. You can either integrate with your CI/CD process to provide the optimal configuration established from prior releases, or when you redeploy we can set your Lambda configuration to this optimal level.



Integrates with your existing tools & workflow

Sedai integrates with popular tools including Monitoring and APM platforms (e.g., Datadog, Prometheus), notification providers (e.g., Slack, PagerDuty), ITSMs (e.g., ServiceNow), and runbook automations (e.g., Rundeck). Sedai also allows you to add additional custom metrics and establish your own SLOs.

About Sedai

Sedai delivers the first autonomous cloud management platform that detects and proactively addresses potential issues in production, improving performance, ensuring availability, and managing cloud costs. Acting as an intelligent autopilot for SREs, Sedai eliminates significant toil for SREs so they can scale and increase innovation cycles. Sedai enhances your Observability and AIOps platforms by proactively preventing issues. Sedai's investors include Norwest Venture Partners and Sierra Ventures.

Interested in learning more?

Email contact@sedai.io and visit www.sedai.io for more information.

