Journal of Artificial Intelligence & Cloud Computing

SCIENTIFIC Research and Community

Review Article

Open 🖯 Access

Optimizing Observability: A Deep Dive into AWS Lambda Metrics

Balasubrahmanya Balakrishna

Senior Lead Software Engineer, Richmond, VA, USA

ABSTRACT

The importance of observability in AWS Lambda systems is examined in this technical article, which focuses on using Python and the AWS Lambda Powertools to improve monitoring capabilities. With serverless computing, AWS Lambda has emerged as a critical service that lets developers concentrate on writing code rather than maintaining infrastructure. Nonetheless, troubleshooting, performance improvement, and guaranteeing the dependability of serverless apps depend on effective observability. Using specially designed Python Serverless Restful API and AWS Lambda Powertools, this article thoroughly analyzes important AWS Lambda metrics, their interpretation, and techniques to improve observability while running serverless applications in the AWS Lambda environment.

The author will use Python with the AWS Lambda Powertools, implemented as a Lambda layer, to facilitate seamless integration for custom metrics and advanced observability features.

*Corresponding author

Balasubrahmanya Balakrishna, Senior Lead Software Engineer, Richmond, VA, USA.

Received: November 07, 2022; Accepted: November 16, 2022; Published: November 25, 2022

Keywords: AWS Lambda, CloudWatch, CloudWatch Metrics, AWS Powertools for Lamda

Background: AWS Cloud Watch Metrics

Two methods are available with AWS Lambda for gathering and examining metrics:

1. Default Option

Without requiring any setup, the default configuration streamlines the monitoring process for users by automatically gathering and presenting critical metrics [1].

2. Custom Option

AWS Lambda allows users to set custom metrics for more customized observability [2]. By allowing the gathering of particular data points, this option makes it possible to monitor and analyze data in a more detailed and unique manner.

Lambda Emits a Default Set of Metrics:

Invocation Metrics

Provides information about the outcome of an Invocation. E.g., Invocations, Errors, Throttles, Dead Letter Errors, Destination Delivery Failures, Provisioned Concurrency Invocations, Provisioned Concurrency Spillover Invocations.

• Dead Letter Errors and Destination Delivery Failures are applicable for Async invocations

Performance Metrics

Provides data about the performance outcome of a single invocation.

E.g., Duration, Post Runtime Extension Duration, Iterator Age, Offset Lag.

• Duration is rounded to the nearest millisecond

- PostRuntimeExtensionDuration is applicable with Lambda extensions
- IteratorAge applies to stream-based
- Offsetlag applies to self-managed Kafka or MSK

Concurrency Metrics

Provides data about the aggregate count of instances processing events.

E.g., Concurrent Executions, Provisioned Concurrent Executions, Provisioned Concurrency Spillover Invocations, Provisioned Concurrency Utilization, Unreserved Concurrent Executions.

- Information can be at the function, version, alias, or region scope
- Unreserved Concurrent Executions represents a region and provides the number of events that function without reserved concurrency are processing
- Concurrent Executions gives the number of function instances
 processing requests
- Provisioned Concurrent Executions gives the number of function instances processing events on provisioned concurrency
- Provisioned Concurrency Utilization gives the value Provisioned Concurrent Executions divided by the amount of provisioned concurrency for a version or alias

Metrics com	e with	predefined	retention	and	resolution	periods,
as shown in	the tab	le below:				

Resolution Period	Retention Period
<1 min*	3 hours
60 sec	15 days
300 sec	63 days
3600 sec	455 days

Citation: Balasubrahmanya Balakrishna (2022) Optimizing Observability: A Deep Dive into AWS Lambda Metrics. Journal of Artificial Intelligence & Cloud Computing. SRC/JAICC-172. DOI: doi.org/10.47363/JAICC/2022(1)160

Initially, data points published at a shorter interval undergo aggregation for prolonged storage. As an illustration, when data is gathered at a 1-minute frequency, it remains accessible at a 1-minute resolution for 15 days. Following this initial period, the data persists but undergoes further aggregation, retrievable only at a 5-minute resolution.

Beyond the 63-day mark, the aggregated data experiences another level of consolidation, becoming accessible with a 1-hour resolution. These evolving resolutions cater to different needs over time, balancing granularity with long-term storage efficiency.

Furthermore, the metrics maintain a retention period of 15 months before entering a rolling-out phase. This structured approach to data retention ensures a balance between historical depth and resource optimization.

The table below provides the appropriate usage of statistics for each metric type:

Metric Type	Statistic to Use	
Invocation Metrics	Sum	
Performance Metrics	Average/Max/Percentile Stats	
Concurrency Metrics	Max	

Introduction

Application Load Balancer (ALB)-fronted Lambda API Architecture

In this exploration, we will spotlight the practical implementation of observability concepts through an Application Load Balancer (ALB)-fronted Lambda API architecture, as depicted in Figure 1. This architectural choice underscores the real-world applicability of the discussed observability strategies.



Figure 1: High-Level Architecture of ALB-Fronted Lambda API

Purpose-Built Lambda Function with AWS Powertools

The key AWS Lambda metrics, their interpretation, and strategies to elevate observability throughout this paper. By utilizing Python and AWS Lambda Powertools, the aim is to empower developers and system administrators with practical insights into building resilient, high-performing, serverless applications. The chosen architectural context and purpose-built function serve as tangible examples of the efficacy of these observability strategies in realworld scenarios.

Optimizing Metric Generation: Leveraging Cloudwatch EMF and Lambda Powertools

Leveraging CloudWatch EMF

Custom metrics can be generated inside a chosen namespace by facilitating the ingestion of logs containing application data using the CloudWatch Embedded Metrics Format (EMF)[3]. These customized metrics are crucially produced asynchronously. A single EMF object can aggregate 100 metrics within a custom namespace. It is not recommended to synchronously report a statistic to CloudWatch Metrics since it hurts function scalability and code performance. Steer clear of such methods.



Figure 2: EMF Output in CloudWatch Logs

It is highly advised for developers working with Python, Typescript, or Java to utilize AWS Lambda Powertools [4]. This toolbox offers a lot of functionality, including analytics, tracing, logging, and other features. Most notably, it makes it possible to generate logs in the proper EMF format quickly.

Exploring Metric Math Capabilities

Although Metric Math is a feature often disregarded, it is more comprehensive than AWS Lambda and Amazon CloudWatch [5]. Once a measure is recorded in CloudWatch, you can use mathematical expressions to create new time series based on the existing metrics, whether custom or built-in. One well-known example in AWS Lambdas is calculating the mistake rate by dividing the Errors metric by the Invocation metric. Metric analysis gains further adaptability by allowing alarms to be created using metrics-related math calculations.

Lambda Function with AWS Powertools Producing Default Metrics on Invocation

The default metrics (Figure 4) are automatically generated as part of the default behavior when the Application Load Balancer (ALB) executes the sample API using AWS Powertools for Python, shown in Figure 3. As the code shows, custom metrics have yet to be expressly defined. Citation: Balasubrahmanya Balakrishna (2022) Optimizing Observability: A Deep Dive into AWS Lambda Metrics. Journal of Artificial Intelligence & Cloud Computing. SRC/JAICC-172. DOI: doi.org/10.47363/JAICC/2022(1)160



Figure 3: Purpose-Built Lambda Function with AWS Powertools

rtitled graph ∠	1h 3h 12	h 1d 3d 1w Custom 🗄	Local timezone V Actions	Number	• 0 •
7.81 ms	16 • frees	177	90.1 s	25 • Concurrent	.6
0.08	0	112 ms	1		
I Throttles	AsyncEventsDropped	AsyncEventAge	AsyncEventsReceived		
		=			
nowse Multi source query - new	Graphed metrics (9) Options Sour	=	inter (multiple) - Z Backet	Add math V	Add query
Add dynamic label	Graphed metrics (9) Options Sour	= ce Stat	istic: (multiple) V Z Peric Statistic I	Add math V ad: 6 hours V	Add query Clear graph Actions
rowse Multi source query - new Add dynamic label V Info	Graphed metrics (9) Options Sour	ee Stat	istic: (multiple) ▼ ∠ Peric Statistic I a∠ Average ▼ 0	Add math ¥ d: 6 hours ¥ Period ¥axis i hours ¥ <>	Add query Clear graph Actions ~ Q (
Add dynamic label V Info	Graphed metrics (9) Options Sour	e Stat	istic: (multiple) ▼ ∠ Perio Statistic I o∠ Average ▼ 0 Iting-tar ∠ Average ▼ 0	Add math V d: 6 hours V Period V axis hours V <> i hours V <>	Add query Clear graph Actions & Q I & Q I
Ironse Multi source query - new Add dynamic labet Label Label Label AnyncivensUropped AnyncivensUropped AnyncivensUropped AnyncivensUropped	Graphed metrics (5) Options Sour Details Lambda - Throttier - Lambda - AsyncEver Lambda - AsyncEver	ce Stat	istic: (multiple) ▼ ∠ Perio Statistic / Average ▼ 0 liling-lar ∠ Average ▼ 0 ambda ∠ Average ▼ 0	Add math V d: 6 hours V Period V axis hours V 5 hours V 6 hours V 6 hours V 6 >	Add query Clear graph Actions Actions Actions

Figure 4: CloudWatch Metrics Generated By Default

Lambda Function with AWS Powertools Produces Custom Metrics on Invocation

Augment the code mentioned above by introducing custom metrics to gain deeper insights into business-level metrics. The modified code snippet below (Figure 5) demonstrates the incorporation of custom metrics. It's important to note that this example merely scratches the surface of the extensive capabilities that Powertools offers for accomplishing a wide range of tasks.

inport	
from per	labbia movertooli immert Lanner, Tracer, Metrica
	landda powertosis.metrics impert PetricUnit
Tron and	Landda powertools.event handler import Aldinoolver, Reiponia, Content_types
from ser	landda powertosla.utilites.typing import LanddaContext
	lands powertools, logging import correlation paths
Instant -	Tracer(service="BooksService")
	 METERSCHER AND AND AND AND AND AND AND AND AND AND
SUCCESS FAILURE	HTTP_STATUS_CODES + [200, 202, 200] HTTP_STATUS_CODES + [401, 400]
SUCC	ESS - Income
Line)	Out - uninger
(tracer.	Capiture, serbad
def get	books() -> Response:
	tatus 👄 Status, SUCCES;
	HTTP_STATUX_CODE = FANDOR_CONCESS_HTTP_STATUS_CODES) THEALT = {
	Tessarce: Lapp.correct.event.pith. Thooks': L'books', "books', "books' I
	<pre>logger.inte["Soccessfully_retrieved_books", extracted()] setrics.add_metric[enney"ResistrifeTercess", uniteRetricBelt.Count, value=0]</pre>
	status_code+http_status_code, content_tenerentent_tenes_allt_trattme_tene
	body=(sch.dumritress(t))
	http:status.cods': http:status.code,
	metrics.add.metric(name="BooksApiFailure", unit=MetricAnit.Count, value=1)
	status code-otto_status_cole,
	content_type=cantent_types_AFFEICATION_3500,
else	TNI
	except Ecception as at an and a set of the s
	logger, exception "Exception occurred", extrav("exception.message': str(s)))
off land	Copilare method
lege	er.infof/'Not_found_route:_{app.current_event.path}*)
	inject lambda context[carrelation_id_pathacerrelation_paths_APM_ICATION_IGAN_BALANCEN_
log_ever	t-True, claur_state-True)
otrace.	Copiure_Lands_handler
def last	da handler(event: dirt, context: LambdaContext) -> Response:

Figure 5: Modified API with Custom Metrics

Within the code, we have established a custom namespace labeled BooksService and populated it with business-level metrics. Figure 6 and Figure 7 below illustrate the specific metric within this custom namespace.

Metrics (262) Info			
N. Virginia 🔻	Q Search for any metric		
Custom namespaces			
BooksService	4		

Figure 6: Custom Namespace



Figure 7: Custom Metrics

Conclusion

AWS Lambda is a critical player in the serverless computing space, freeing developers from the burden of managing infrastructure in favor of writing code. This exploration of observability optimization in AWS Lambda environments has revealed a variety of tactics meant to improve accuracy and efficiency. Exploring default metrics categories like Performance, Concurrency, and Invocation has helped us establish a solid foundation for debugging and monitoring.

Furthermore, a potentially has appeared with the advent of AWS Lambda Powertools, which are smoothly linked with Python. One helpful step toward a more sophisticated observability framework is using the CloudWatch Embedded Metrics Format (EMF) and the advice to use Powertools for metric creation, tracing, and logging.

Beyond the technical details, investigating Metric Math capabilities and adding custom metrics under a namespace designated explicitly for them (such as "BooksService") highlight the breadth of analysis possible with AWS Lambda and CloudWatch.

As we go through the observability's complexities, it is clear that this journey is about more than just measurements and codes-it's about making decisions based on intelligent facts. Refinements for optimal system knowledge and responsiveness are prioritized, whether avoiding anti-patterns like synchronous metric publication or utilizing Metric Math for meaningful computations.

Essentially, the cooperative synergy of AWS Lambda, CloudWatch, Python, and Powertools provides a comprehensive approach to observability, which guarantees the identification of problems and proactive optimization for continued performance. Developers and system administrators can create a serverless environment that functions effectively and changes dynamically in response to the needs of the applications it supports by following best practices, avoiding hazards, and utilizing the capabilities offered. Pursuing increased observability is a process rather than a destination, and this investigation is the first step toward a day when AWS Lambda applications function at their best.

References

- Metrics collected by Lambda Insights. Amazon CloudWatch User Guide https://docs.aws.amazon.com/ AmazonCloudWatch/latest/monitoring/Lambda-Insightsmetrics.html.
- 2. Publish custom metrics. Amazon CloudWatch User Guide https://docs.aws.amazon.com/AmazonCloudWatch/latest/

monitoring/publishingMetrics.html.

- Setting alarms on metrics created with the embedded metric format. Amazon CloudWatch User Guide https://docs. aws.amazon.com/AmazonCloudWatch/latest/monitoring/ CloudWatch_Embedded_Metric_Format_Alarms.html
- Metrics. AWS Powertools for AWS Lambda (Python) https:// docs.powertools.aws.dev/lambda/python/latest/core/metrics/.
- 5. Use metric math. Amazon CloudWatch User Guide https://docs.aws.amazon.com/AmazonCloudWatch/latest/ monitoring/using-metric-math.html.

Copyright: ©2022 Balasubrahmanya Balakrishna. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.