

Serverless updates from
AWS re:Invent 2025



Show of hands





The biggest secret of Serverless



(and many other serverless services)



Hundreds of thousands of
EC2 instances under-the-hood



© 2025, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Usman Khalid

He/His/Him

Head of Serverless Compute

linkedin.com/in/ahmedusmankhalid/

hundreds of thousands of servers.



What actually is serverless

Managed by AWS

- management of servers
- management of runtimes
- scaling
- load balancing
- request routing
- patching
- ...

Pricing model

- pay per request
- pay per usage
- pay per duration
- scaling to 0
- on demand pricing

“I don’t want to manage my own cluster - that’s what AWS is for.”

AWS serverless

Today's focus:

- AWS Lambda
- AWS ECS
- DynamoDB & DSQL



a full list of
AWS serverless services

AWS Lambda

- FaaS (Function as a Service)
- various runtimes (official and unofficial)
- Docker lambda, Lambda Web Adapter
- **128MB - 10 240MB RAM, 1 - 6 vCPU**
- **15 minutes timeout**
- rapid auto scaling
- fully managed
- Event Source Mapping (ESM)



Lambda Managed Instances (LMI)

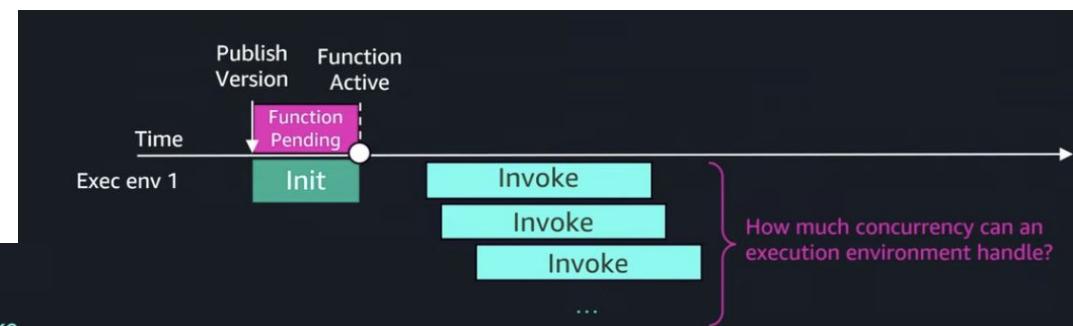
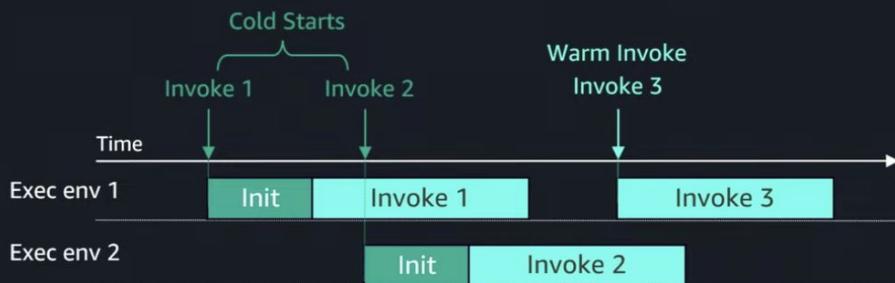
- pick underlying EC2 instance type
- fully managed
- same model as Lambda (no need to refactor/re-architect)
- great for steady and CPU heavy workloads (for unpredictable use standard Lambda)
- multi-concurrency with less cold starts
- you pay for EC2 + 15% management overhead
- supports Saving plans and Reserved Instances

address compute choice, cost and cold starts limitations

Lambda Managed Instances (LMI)

1. create capacity provider (optionally select instance types)
2. associate your lambda function with capacity provider (adjust memory/cpu ratio)
3. deploy your function

Standard Lambda 



LMI multi-concurrency 

```
1 import { App, Stack } from "aws-cdk-lib";
2 import { SecurityGroup, Vpc } from "aws-cdk-lib/aws-ec2";
3 import {
4     CapacityProvider,
5     ScalingOptions,
6     TargetTrackingScalingPolicy,
7 } from "aws-cdk-lib/aws-lambda";
8
9 const stack = new Stack(new App(), "MyStack");
10 const vpc = new Vpc(stack, "MyVpc");
11 const securityGroup = new SecurityGroup(stack, "SecurityGroup", { vpc });
12
13 const capacityProvider = new CapacityProvider(stack, "MyCapacityProvider", {
14     subnets: vpc.privateSubnets,
15     securityGroups: [securityGroup],
16     scalingOptions: ScalingOptions.manual([
17         TargetTrackingScalingPolicy.cpuUtilization(70),
18     ]),
19 }
20 ) architectures? (property) CapacityProviderProps.architectures?
21 capacityProviderName? (property) CapacityProviderProps.capacityProviderName?
22 instanceTypeFilter? (property) CapacityProviderProps.instanceTypeFilter?
23 kmsKey? (property) CapacityProviderProps.kmsKey?: IKey | undefined
24 maxVCpuCount? (property) CapacityProviderProps.maxVCpuCount?: number
25 operatorRole? (property) CapacityProviderProps.operatorRole?: string
```

The instruction set architecture required for compute instances.
Only one architecture can be specified per capacity provider.

`@default`

- No architecture constraints specified

```
24
25 const lambda = new NodejsFunction(stack, "MyLambda", {
26   entry: "handler.js",
27   runtime: Runtime.NODEJS_24_X,
28   architecture: Architecture.ARM_64,
29   memorySize: 1024,
30 });
31 capacityProvider.addFunction(lambda, {
32   executionEnvironmentMemoryGiBPerVCpu: 4,
33   perExecutionEnvironmentMaxConcurrency: 20,
34
35 }) latestPublishedScalingConfig? (property) CapacityProviderFunctionC
36 publishToLatestPublished? (property) CapacityProviderFunctionC
```

The scaling options that are applied to the \$LATEST.PUBLISHED version.

@default

- No scaling limitations are applied to the \$LATEST.PUBLISHED version.

Lambda: Durable Functions

- workflow orchestration within single AWS Lambda
- powerful SDK (Node.js and Python, more are coming soon)
- you don't pay for waiting (up to 1 year execution time)
- checkpoints and replays
- similar to Step Functions

address timeout and orchestration limitations

```
1 import { Duration } from ' ../../core';
2 /**
3  * Configuration for durable functions.
4  *
5  * Lambda durable functions allow for long-running executions with persistent state.
6 */
7 export interface DurableConfig {
8     /**
9      * The amount of time that Lambda allows a durable function to run before stopping it.
10     *
11     * Must be between 1 and 31,622,400 seconds (366 days).
12     */
13     readonly executionTimeout: Duration;
14     /**
15      * The number of days after a durable execution is closed that Lambda retains its history.
16     *
17     * Must be between 1 and 90 days.
18     *
19     * The underlying configuration is expressed in whole numbers of days. Providing a Duration that
20     * does not represent a whole number of days will result in a runtime or deployment error.
21     *
22     * @default Duration.days(14)
23     */
24     readonly retentionPeriod?: Duration;
25 }
```

Durable Functions demo - what I've learned

1. problems with permissions
2. old SDK in Node.js runtime
3. permissions setup



public gist
with demo code

Durable operations (3) Info										Last fetched 30 seconds ago C		
	Name	Status	SubType	Timeline	Start	End	Duration	Re...	Actions			
○	wait-for-extern...	✓ Succeeded	WaitForCa...	<div><div style="width: 100%;"> </div></div>	Jan 27, 2026 at 14:45:19.309 (UTC+01:00)	Jan 27, 2026 at 14:45:22.316 (UTC+01:00)	3 secs 7 ms	-	-			
○	ea66c06c1...	✓ Succeeded	Callback	<div><div style="width: 50%;"> </div></div>	Jan 27, 2026 at 14:45:19.309 (UTC+01:00)	Jan 27, 2026 at 14:45:21.044 (UTC+01:00)	1 sec 735 ms	-	-			
○	98c6f2c228...	✓ Succeeded	Step	<div><div style="width: 50%;"> </div></div>	Jan 27, 2026 at 14:45:19.377 (UTC+01:00)	Jan 27, 2026 at 14:45:21.106 (UTC+01:00)	1 sec 729 ms	-	-			

Step-Functions vs. Durable functions

Choosing a service

Use to orchestrate across AWS services

If you want a visual builder experience

Never want to manage or patch

Use when orchestrating at the application level

If you want to use your favorite programming language

Don't mind updating your functions when needed



Step-Functions vs. Durable functions - Eric Johnson



CNS380
Durable Functions



API313
StepFunctions

AWS Lambda - worth mentioning

- public Lambda Roadmap on Github
- Tenant Isolation (pre:Invent)
- support for Egress-only Internet Gateway (IPv6 only)
- StepFunctions JSONata support (lightweight transformations without Lambda)

The screenshot shows the AWS Lambda roadmap on GitHub, organized into five columns:

- Researching** (40 items):
 - aws-lambda-roadmap #19: Enable developers to augment AWS SAM and SAM CLI capabilities ("AWS SAM Plugins")
 - aws-lambda-roadmap #20: OpenTelemetry (OTel) support
 - aws-lambda-roadmap #23: AWS SAM support for API Gateway WebSocket
 - aws-lambda-roadmap #29: [Redacted]
- Working On It** (6 items):
 - aws-lambda-roadmap #30: Flexibility to set block function create and update date for deprecated Lambda managed runtimes
 - aws-lambda-roadmap #32: JSON-based Resource-based Policies
 - aws-lambda-roadmap #33: Block public access for functions [Resource level]
- Coming Soon** (4 items):
 - aws-lambda-roadmap #18: Recursive loop detection and recursive loop detection APIs in all commercial AWS Regions
 - aws-lambda-roadmap #49: Enhanced observability (CloudWatch Logs and metrics) for Kafka event source mappings (ESMs)
 - aws-lambda-roadmap #52: Cross-account access for DynamoDB Streams
- Developer Preview** (0 items):

This item is in preview
- Shipped** (21 items):
 - aws-lambda-roadmap #27: Lambda managed runtime for .NET 10
 - aws-lambda-roadmap #15: Lambda managed runtime for Node.js 24
 - aws-lambda-roadmap #4: Provisioned Mode for Kafka event source mappings (ESMs)
 - aws-lambda-roadmap #1: [Redacted]

Elastic Container Service (ECS)

- native container orchestration service
- **compute service - Fargate (serverless), EC2**
- no control plane to manage
- blue/green deployments
- services and tasks
- pay as go model for used resources
- used for long running (15m+) serverless tasks before Durable Functions
- various deployment options (CDK, CloudFormation, AWS Copilot)



ECS Managed Instances (CNS342)

- simplicity of Fargate with control of EC2
- optional selection of instance types
- instances deployed to customer account
- fully managed by AWS
- ~2w retention lifecycle of instances
- cost optimized
- new capacity provider for EC2 Managed Instance

address Fargate flexibility

DynamoDB

- schemaless NoSQL key-value store
- highly available, scalable,
fully managed, serverless storage
- strong consistency and eventual consistency modes
(both regional and multi-regional)



“With DynamoDB, you are the query planner.”
- Alex DeBrie at DynamoDB Day 2025

DynamoDB - multi-attribute composite keys

- global secondary index (GSI) support
- up to 8 attributes in key
- up to 20 GSI per table (+ 5 LSI)
- no more “country#state#province#city” nonsense
- still works in hierarchy (eg. you need to query from left to right without skipping)
- no table primary key support yet

Aurora DSQL

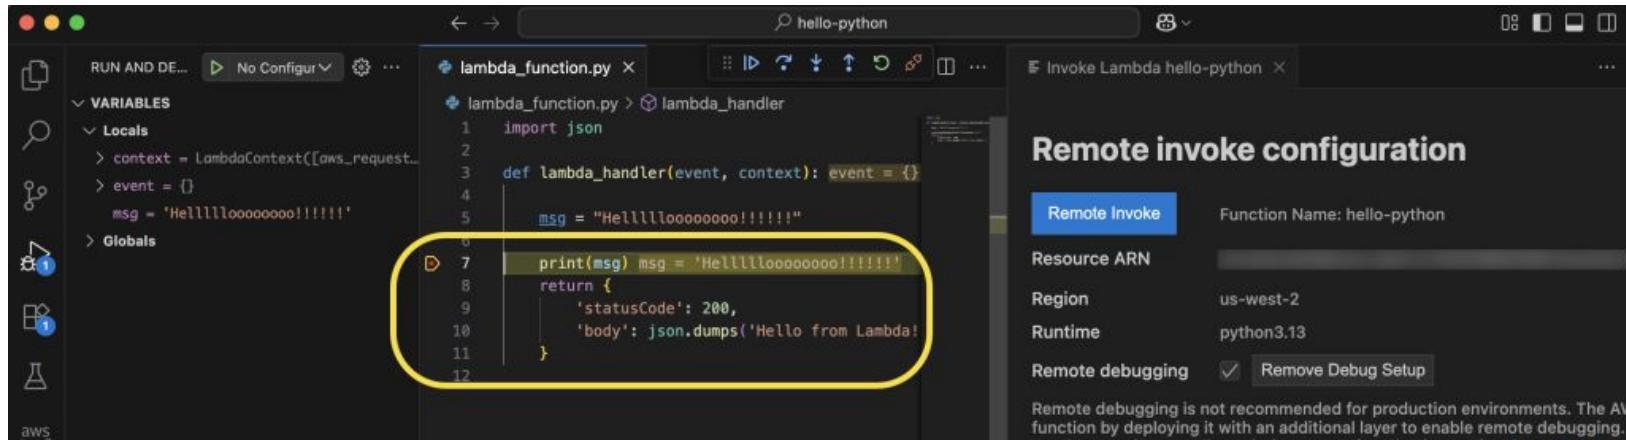
- **distributed** SQL database (postgres compatible)
- fully managed, scalable, serverless (super simple to create and use)
- for those who love their SQL
- no foreign keys constraints support
- cross-region strong consistency by default (everything runs in transaction)
- optimistic concurrency control (OCC) (DAT455 - A Tale of Two Transactions)



if you have time watch deep dive on architecture

Local development

- AWS toolkit for Visual Studio Code
- LocalStack support (dockerized AWS simulator)
- remote debugging support
- Console to IDE (start in console move to VSCode)



The screenshot shows the AWS Toolkit for Visual Studio Code interface. On the left, the sidebar displays 'RUN AND DE...' (No Configur), 'VARIABLES' (Locals and Globals), and 'aws'. The main area shows a code editor with 'lambda_function.py' open, containing the following Python code:

```
lambda_function.py x
lambda_function.py > lambda_handler
1 import json
2
3 def lambda_handler(event, context): event = {}
4
5     msg = "Helllllooooooo!!!!!!"
6
7     print(msg) msg = "Helllllooooooo!!!!!!"
8
9     return {
10         'statusCode': 200,
11         'body': json.dumps('Hello from Lambda!')
12     }
```

A yellow box highlights the line `print(msg) msg = "Helllllooooooo!!!!!!"`. To the right, a panel titled 'Remote invoke configuration' shows the following settings:

Remote Invoke	Function Name: hello-python
Resource ARN	[REDACTED]
Region	us-west-2
Runtime	python3.13
Remote debugging	<input checked="" type="checkbox"/> Remove Debug Setup

At the bottom, a note states: 'Remote debugging is not recommended for production environments. The AW function by deploying it with an additional layer to enable remote debugging.'

re:Watch **unofficial** - AWS re:Invent Video Library



big thanks to
Martin Damovsky



re:Watch playlist
Serverless

Ivan Barlog

AWS Solutions Architect



Github [ivanbarlog](https://github.com/ivanbarlog)

Email ivan@barlog.sk

Web barlog.sk
beesolve.com

LinkedIn [ivan-barlog](https://www.linkedin.com/in/ivan-barlog)

Go build
something!