



Adopting Java for the Serverless world

from the perspective of the **AWS** developer



Contact

Vadym Kazulkin, ip.labs GmbH



v.kazulkin@gmail.com



<https://www.linkedin.com/in/vadymkazulkin/>



@VKazulkin

Co-Organizer of Java User Group Bonn &
Serverless Bonn Meetup

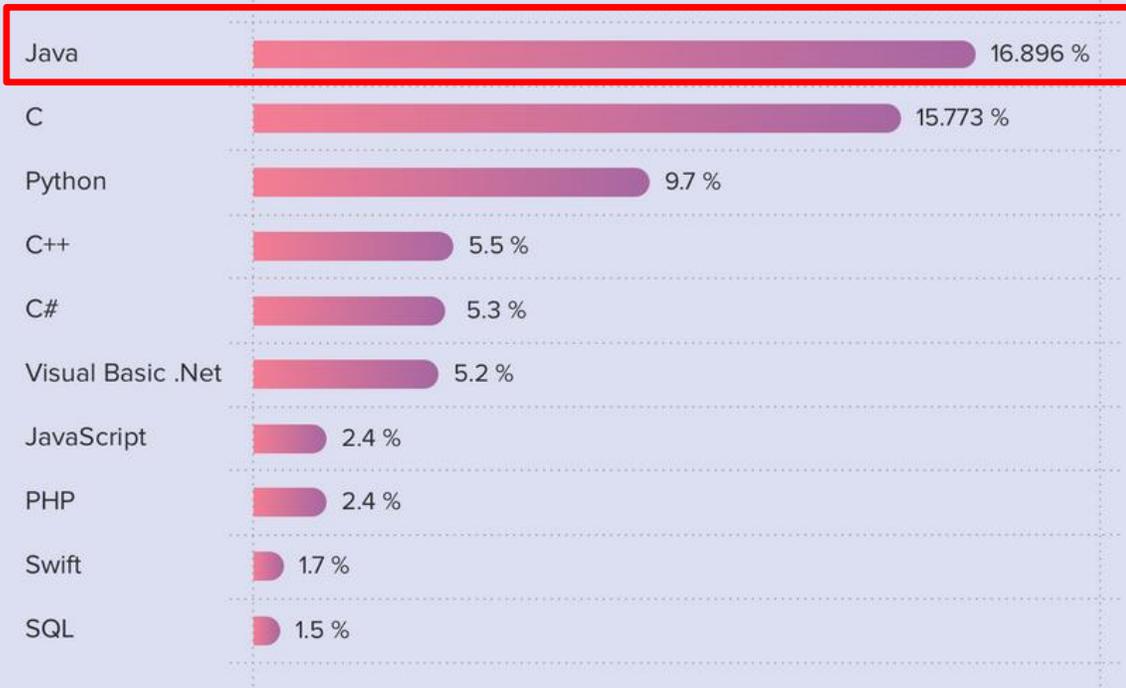






Java popularity

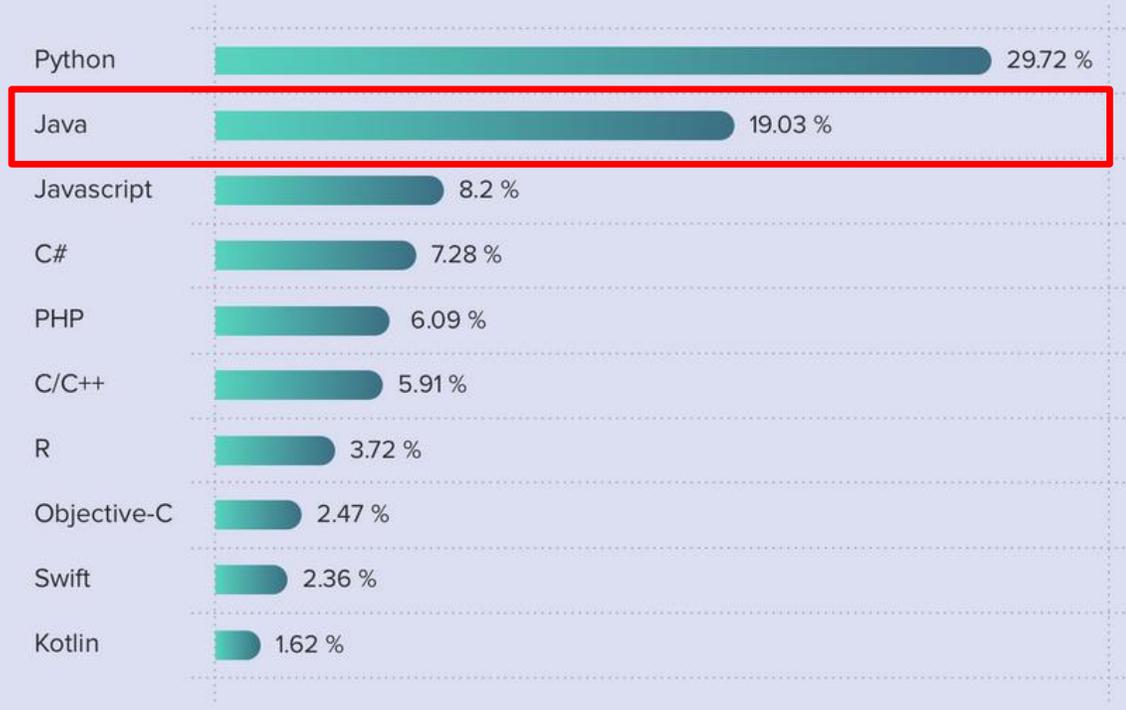
Top programming languages, TIOBE



SHARE

Programming language popularity, 2020 by [TIOBE](#)

Top programming languages, PYPL



SHARE

Most popular coding languages, 2020 by [PYPL](#)

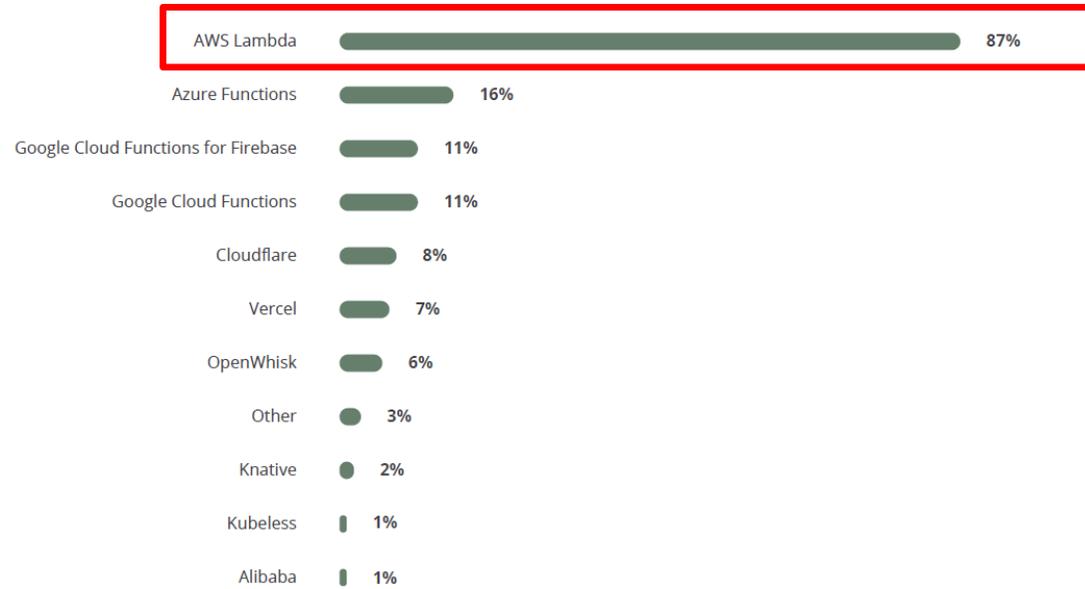


AWS and Serverless

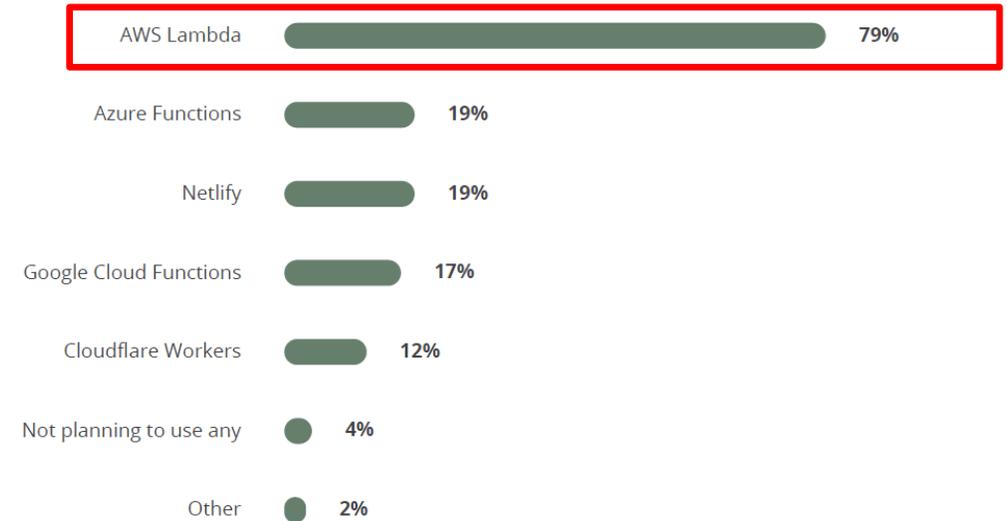
Figure 1. Magic Quadrant for Cloud Infrastructure and Platform Services



Which serverless vendor do you use?

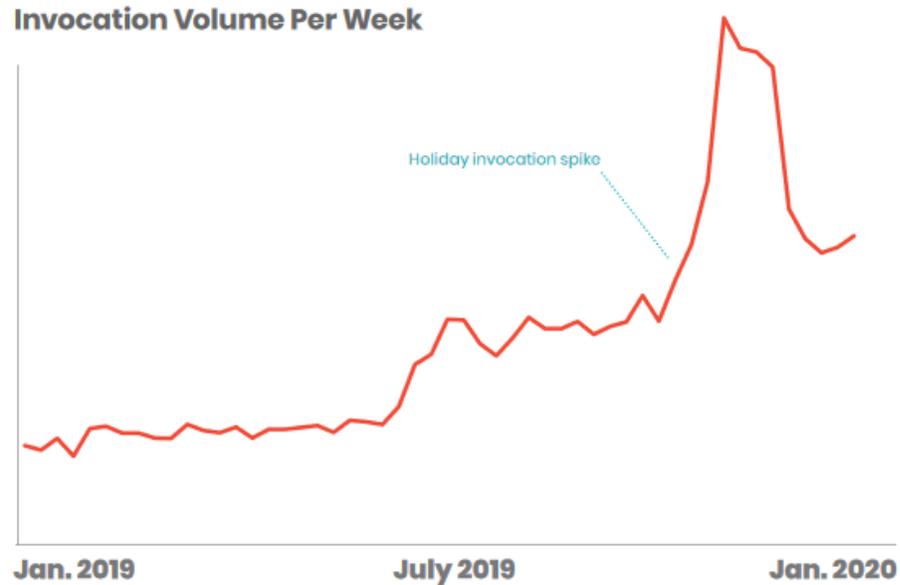


Which FaaS products are you planning to use in the next 12 months?



The Rise in Serverless

Invocation Volume Per Week



Invocation Growth
+ 209%

Average weekly invocation volume within sample set over the past 12 months.

Insights Snapshot

- Serverless adoption continues to rise, with a 209% increase in average weekly invocations over the last 12 months.
- Large spikes in invocations during the holiday season reinforce the serverless use case of auto-scaling to support peak holiday workloads across industries like retail, media, and logistics.



Life of the Java Serverless developer on AWS

Java Versions Support

- Java 8
 - with long-term support
- Java 11 (since 2019)

Amazon Corretto

No-cost, multiplatform, production-ready distribution of OpenJDK

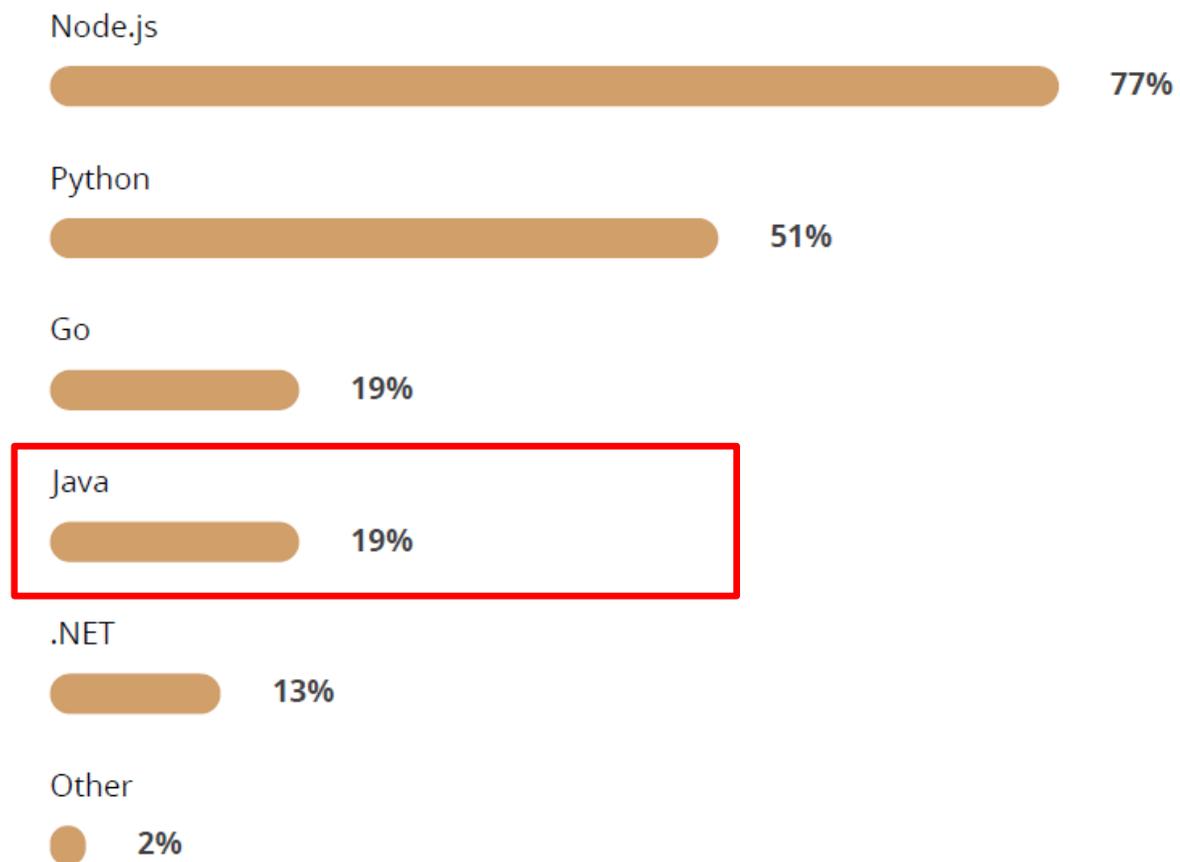
Amazon Corretto is a no-cost, multiplatform, production-ready distribution of the Open Java Development Kit (OpenJDK). Corretto comes with long-term support that will include performance enhancements and security fixes. Amazon runs Corretto internally on thousands of production services and Corretto is certified as compatible with the Java SE standard. With Corretto, you can develop and run Java applications on popular operating systems, including Linux, Windows, and macOS.

Java ist very fast
and mature
programming
language...

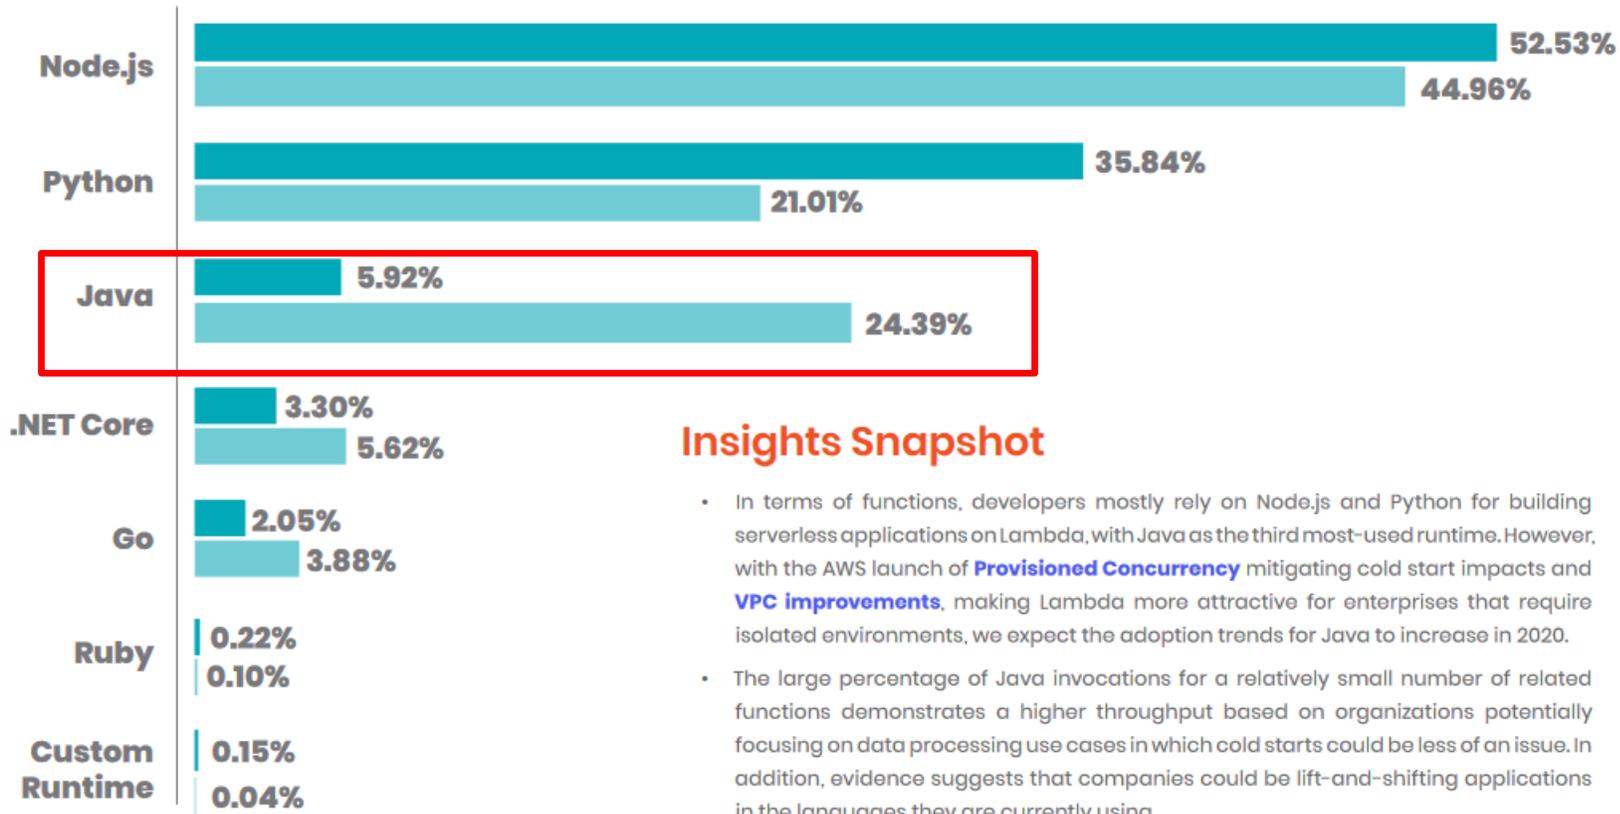
... but
Serverless
adoption of Java
looks like this



What programming languages does your organization use to develop serverless functions?



Lambda Adoption by Runtime



Insights Snapshot

- In terms of functions, developers mostly rely on Node.js and Python for building serverless applications on Lambda, with Java as the third most-used runtime. However, with the AWS launch of **Provisioned Concurrency** mitigating cold start impacts and **VPC improvements**, making Lambda more attractive for enterprises that require isolated environments, we expect the adoption trends for Java to increase in 2020.
- The large percentage of Java invocations for a relatively small number of related functions demonstrates a higher throughput based on organizations potentially focusing on data processing use cases in which cold starts could be less of an issue. In addition, evidence suggests that companies could be lift-and-shifting applications in the languages they are currently using.

Percentage of all **functions** monitored
Percentage of all **invocations** monitored

Time frame July-December 2019



Developers love Java and will be happy
to use it for Serverless

```
C:\Windows\system32\cmd.exe
E:\Java>javac First.java
E:\Java>java First
Let's do something using Java technology.
E:\Java>
```



Creating AWS Lambda with Java 1/2

Author from scratch Start with a simple "hello world" example.

Blueprints Choose a preconfigured template as a starting point for your Lambda function.

Serverless Application Repository Find and deploy serverless apps published by developers, companies, and partners on AWS.

Author from scratch [Info](#)

Name

Runtime

Role
Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.

Existing role
You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.

Basic settings

Description

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.
 MB

Timeout [Info](#)
 min sec

Creating AWS Lambda with Java 2/2

```
import javax.inject.Inject;

import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

import com.amazonaws.services.lambda.runtime.Context;
import com.amazonaws.services.lambda.runtime.RequestHandler;

public class MonthlyInvoiceGeneratorFunction
implements RequestHandler<MonthlyInvoiceRequest, MonthlyInvoiceResponse> {

    private static final Logger LOG = LoggerFactory.getLogger(MonthlyInvoiceGeneratorFunction.class);

    @Inject
    private MonthlyInvoiceGeneratorService monthlyInvoiceGeneratorService;

    @Override
    public MonthlyInvoiceResponse handleRequest(MonthlyInvoiceRequest monthlyInvoiceRequest,
        final Context context) {
        if (LOG.isDebugEnabled()) {
            LOG.debug("request: {}", monthlyInvoiceRequest);
        }
        return this.monthlyInvoiceGeneratorService.generateInvoice(monthlyInvoiceRequest);
    }
}
```

```
import java.util.function.Function;

import org.slf4j.Logger;
import org.slf4j.LoggerFactory;
import org.springframework.beans.factory.annotation.Autowired;

public class MonthlyInvoiceGeneratorFunction
implements Function<MonthlyInvoiceRequest, MonthlyInvoiceResponse> {

    private static final Logger LOG = LoggerFactory.getLogger(MonthlyInvoiceGeneratorFunction.class);

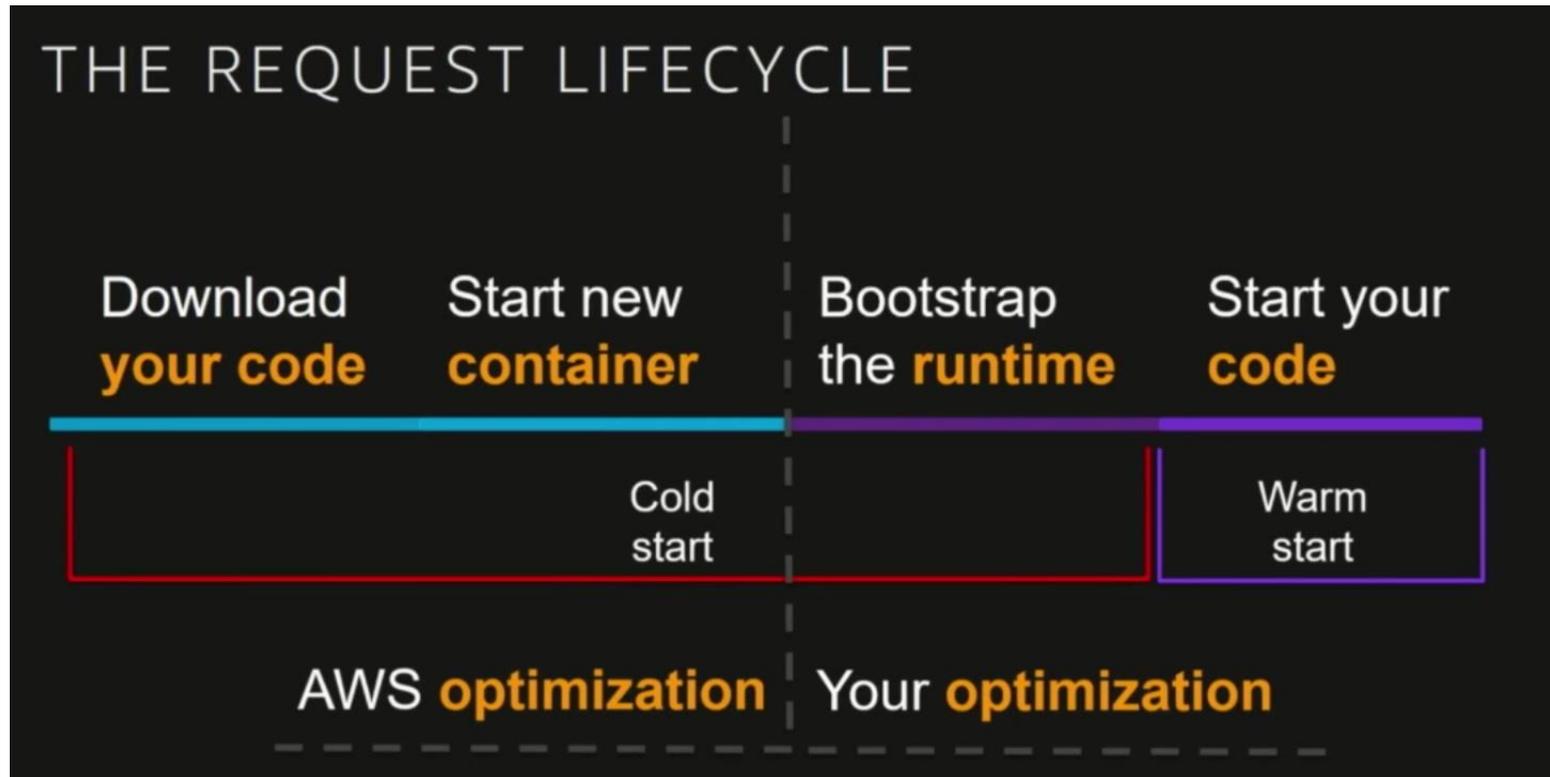
    @Autowired
    private MonthlyInvoiceGeneratorService monthlyInvoiceGeneratorService;

    @Override
    public MonthlyInvoiceResponse apply(MonthlyInvoiceRequest monthlyInvoiceRequest) {
        if (LOG.isDebugEnabled()) {
            LOG.debug("request: {}", monthlyInvoiceRequest);
        }
        return this.monthlyInvoiceGeneratorService.generateInvoice(monthlyInvoiceRequest);
    }
}
```

Challenge Number 1 with Java is a
big **cold-start**

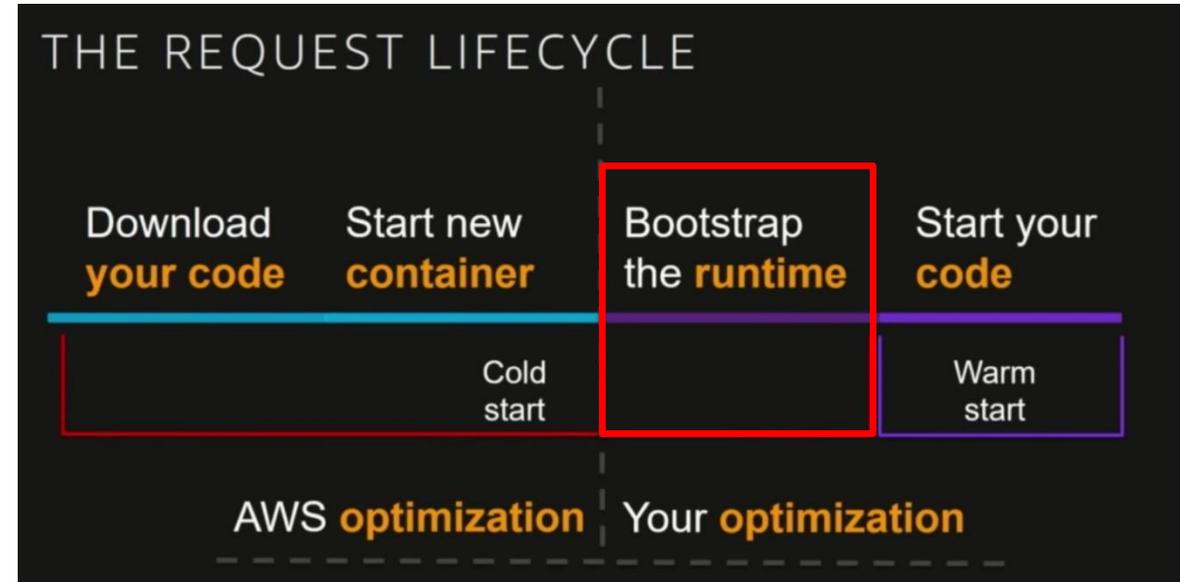


Cold Start



Bootstrap the Java Runtime Phase

- AWS Lambda starts the JVM
- Java runtime loads and initializes handler class
 - Static initializer block of the handler class is executed
 - Boosted host full CPU access up to 10 seconds
- Lambda calls the handler method
 - Full CPU access only approx. from 1.8 GB “assigned” memory to the function



Basic settings

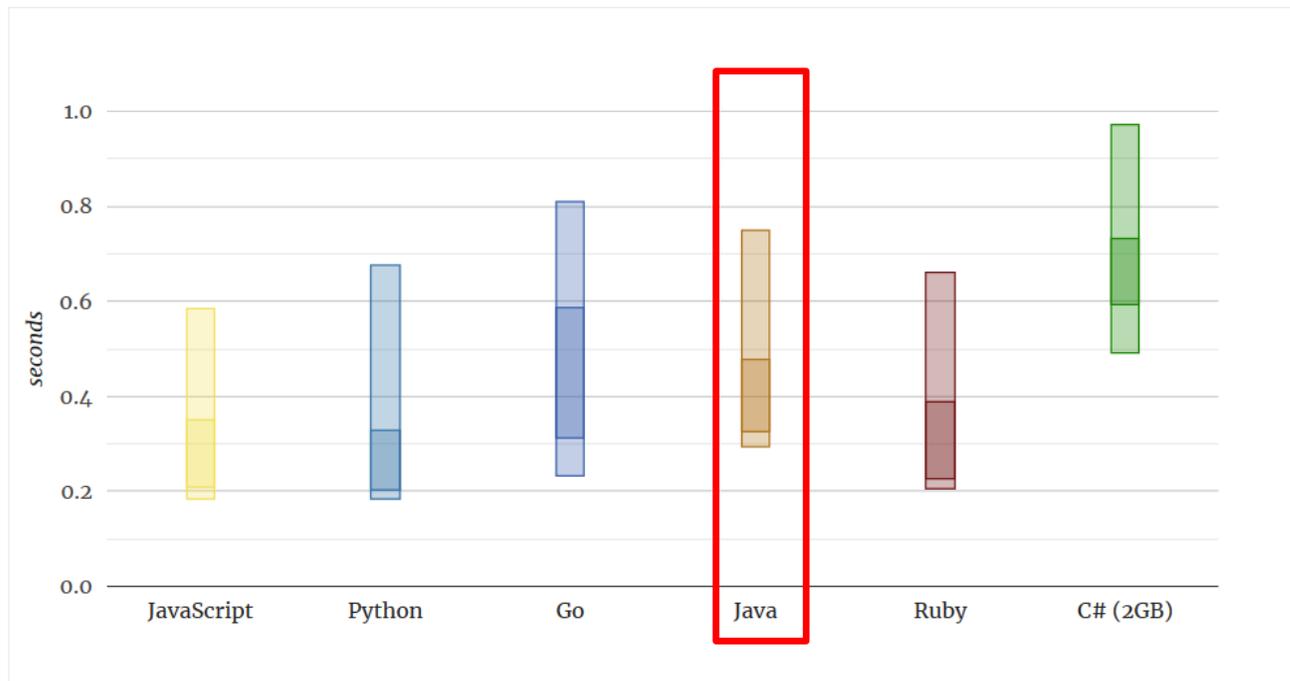
Description

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.
128 MB

Timeout [Info](#)
0 min 3 sec

AWS Lambda cold start duration per programming language

The following chart shows the typical range of cold starts in AWS Lambda, broken down per language. The darker ranges are the most common 67% of durations, and lighter ranges include 95%.

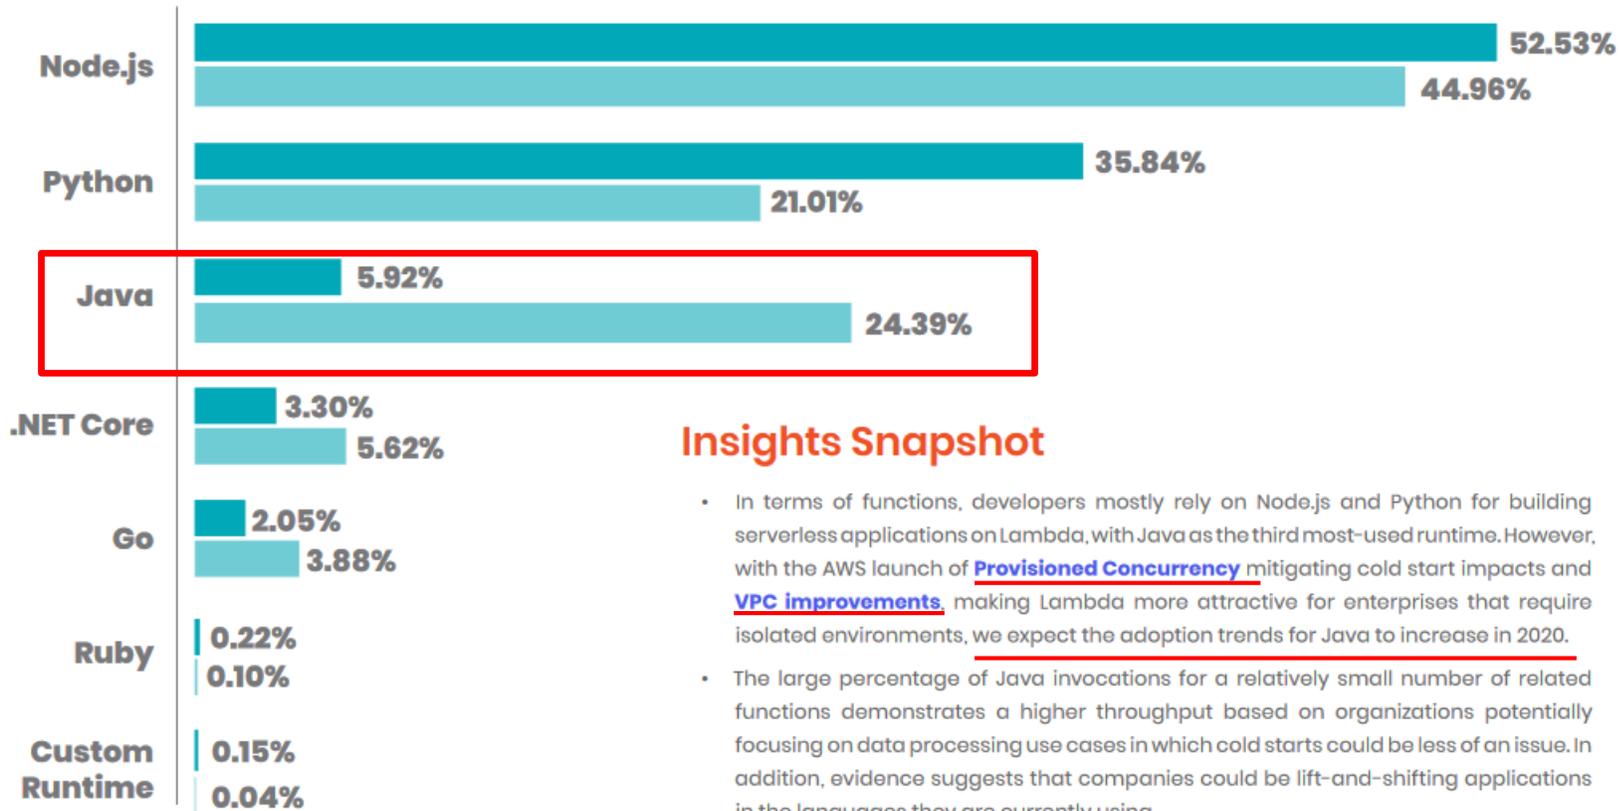


Typical cold start durations per language

Cold start duration with Java

- Below 1 second is best-case cold start duration for very simple Lambda like HelloWorld
- It goes up significantly with more complex scenarios
 - Dependencies to multiple OS projects
 - Clients instantiation to communicate with other (AWS) services (e.g. DynamoDB, SNS, SQS, 3rd party)
- To achieve the minimal cold start duration apply all best practices from these talk
 - Worst-case cold starts can be higher than 10 and even 20 seconds

Lambda Adoption by Runtime



Insights Snapshot

- In terms of functions, developers mostly rely on Node.js and Python for building serverless applications on Lambda, with Java as the third most-used runtime. However, with the AWS launch of [Provisioned Concurrency](#) mitigating cold start impacts and [VPC improvements](#), making Lambda more attractive for enterprises that require isolated environments, we expect the adoption trends for Java to increase in 2020.
- The large percentage of Java invocations for a relatively small number of related functions demonstrates a higher throughput based on organizations potentially focusing on data processing use cases in which cold starts could be less of an issue. In addition, evidence suggests that companies could be lift-and-shifting applications in the languages they are currently using.

Percentage of all **functions** monitored
Percentage of all **invocations** monitored

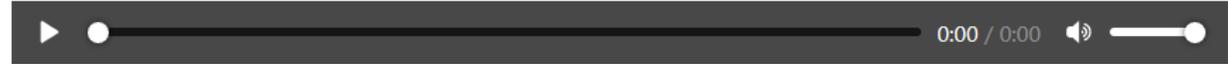
Time frame July-December 2019



Provisioned Concurrency for Lambda Functions

New – Provisioned Concurrency for Lambda Functions

by Danilo Poccia | on 03 DEC 2019 | in [Auto Scaling](#), [AWS Lambda](#), [AWS Re:Invent](#), [AWS Serverless Application Model](#), [Compute](#), [Launch, News](#), [Serverless](#) | [Permalink](#) | [Share](#)

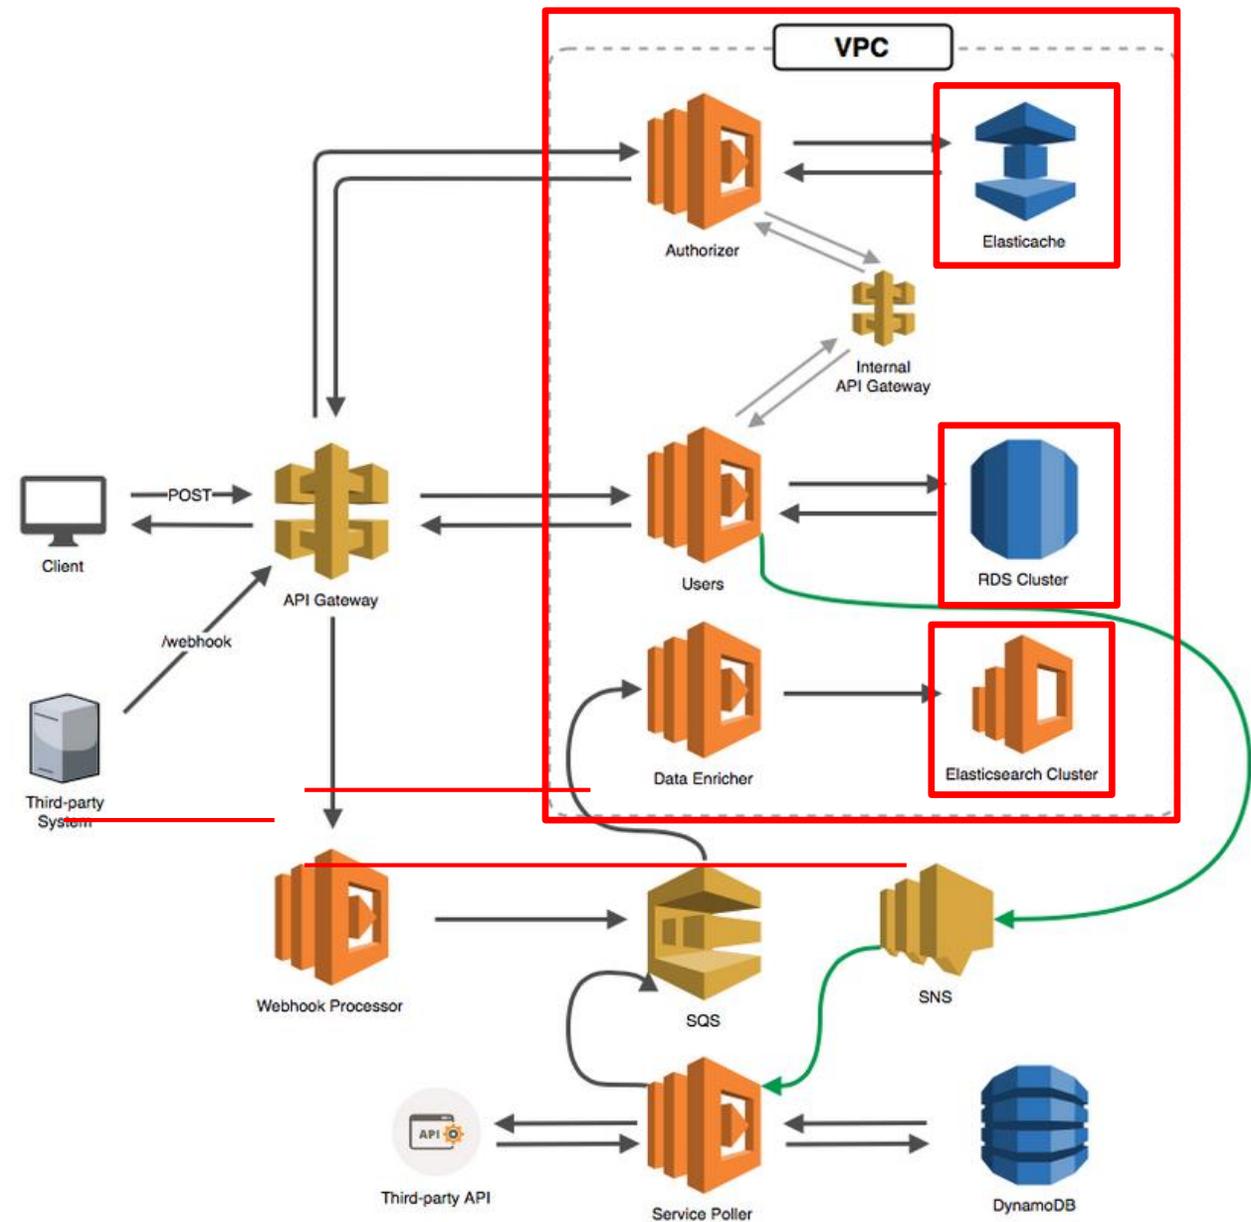


Voiced by [Amazon Polly](#)

It's really true that time flies, especially when you don't have to think about servers: [AWS Lambda just turned 5 years old](#) and the team is always looking for new ways to help customers build and run applications in an easier way.

As more mission critical applications move to [serverless](#), customers need more control over the performance of their applications. Today we are launching **Provisioned Concurrency**, a feature that keeps functions initialized and hyper-ready to respond in double-digit milliseconds. This is ideal for implementing interactive services, such as web and mobile backends, latency-sensitive microservices, or synchronous APIs.

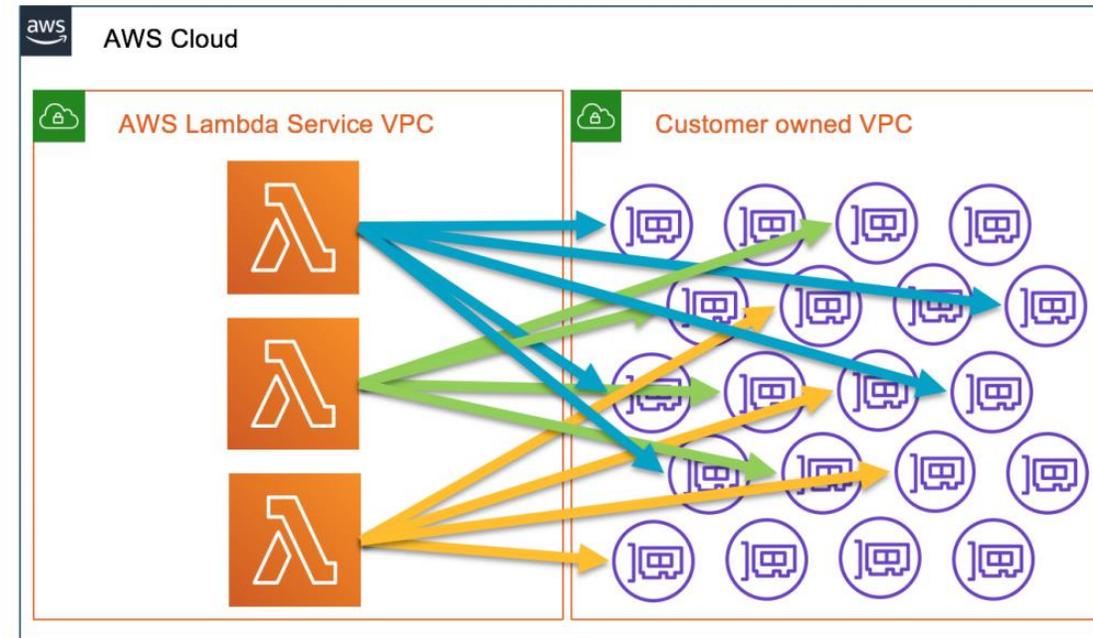
Lambda behind the Virtual Private Cloud (VPC)



Lambda in VPC

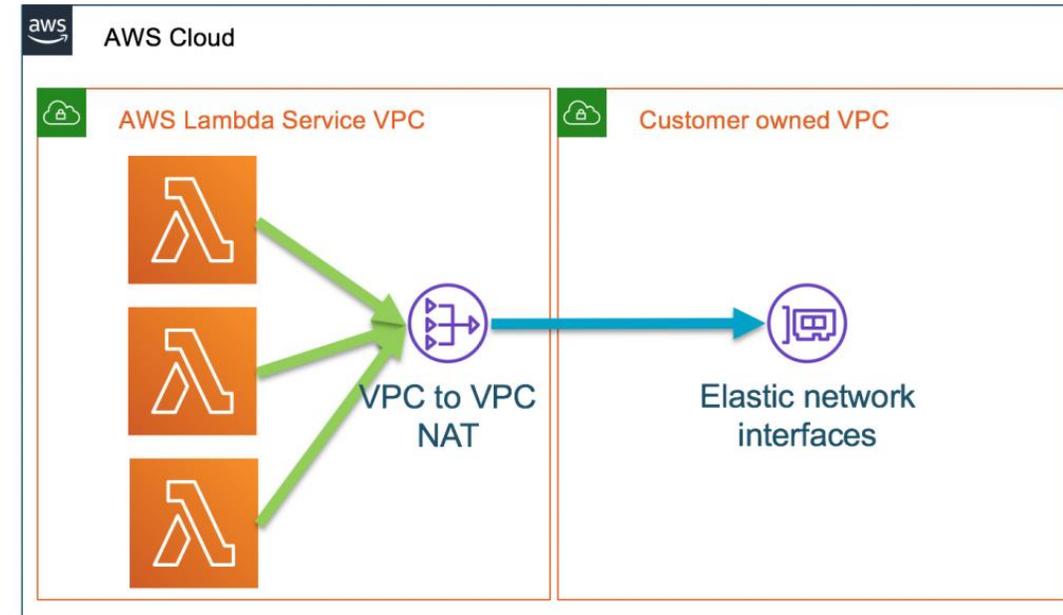
As function's execution environment scales

- More network interfaces are created and attached to the Lambda infrastructure
- The exact number of network interfaces created and attached is a factor of your function configuration and concurrency
- Caused additional the cold start up to approx. **10 seconds**



Lambda in VPC Improvements:

- The network interface creation happens when Lambda function is created or its VPC settings are updated.
- Because the network interfaces are shared across execution environments, only a handful of network interfaces are required per function
- Reduced additional cold start **from approx. 10 seconds to below 1 second**

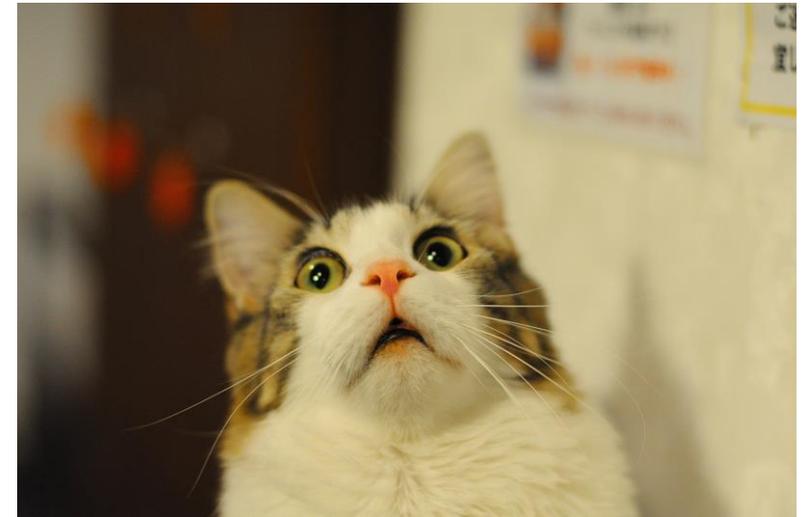


Improvements 1/4

- Switch to the AWS SDK 2.0 for Java
 - Lower footprint and more modular
 - Allows to configure HTTP Client of our choice (e.g. Java own Basic HTTP Client)
- Initialize and prime dependencies during initialization phase
 - Use static initialization in the handler class
- Provide all known values (for building clients e.g. DynamoDBClient) to avoid auto-discovery
 - credential provider, region, endpoint
- Less (dependencies, classes) is more

Improvements 2/4

Avoid Reflection



Or use DI Frameworks like Dagger which aren't reflection-based

Improvements 3/4

Strive for cost optimization



AWS Lambda pricing model

Cost for Lambda



REQUEST



DURATION



Request Tier

\$ 0.20

Per 1 Mio Requests



Duration Tier

\$ 0.00001667

Per GB-Second

GB-Second



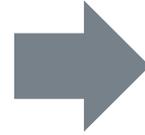
ONE SECOND



ONE GB

Example

- 1 Mio requests
- Lambda with 512MiB
- Each lambda takes 200ms



$0.5 \text{ GiB} * 0.2 \text{ sec} * 1 \text{ Mio}$
 $= 100\,000 \text{ GB-Seconds}$

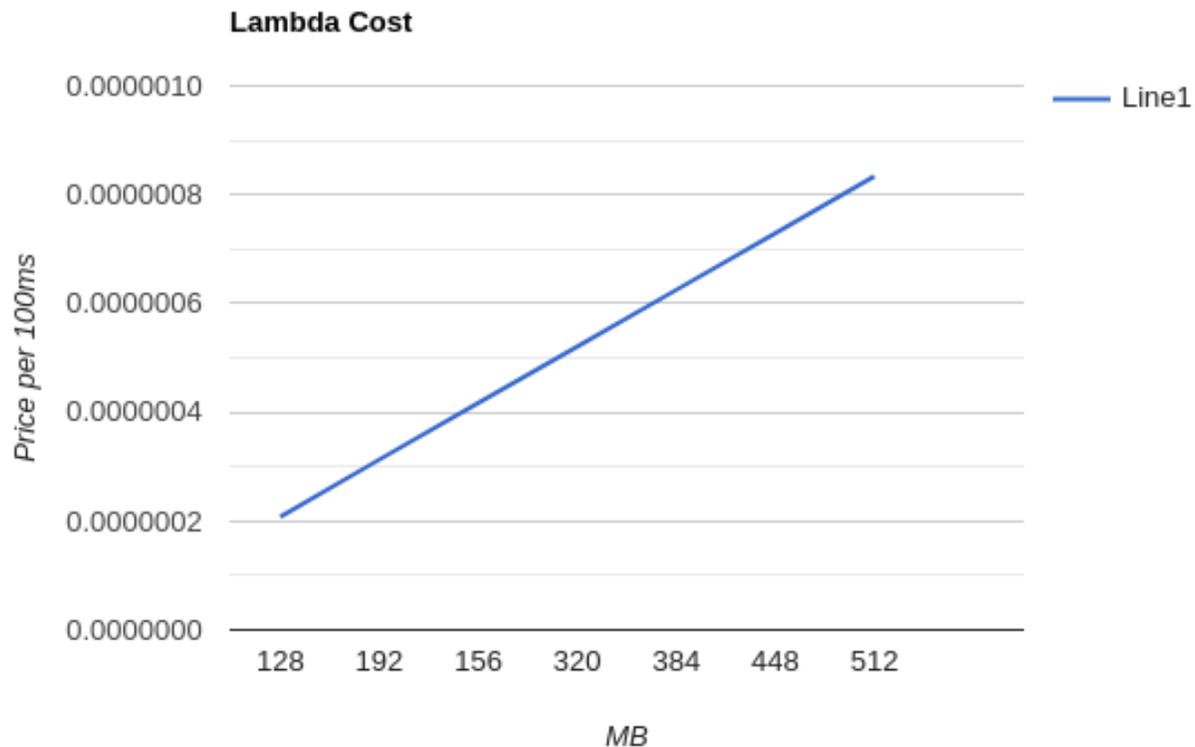


Requests:
\$0.20



GB-Seconds:
\$1.67

Cost scales linearly with memory



Improvements 4/4

More memory = more expensive?

Basic settings

Description

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.

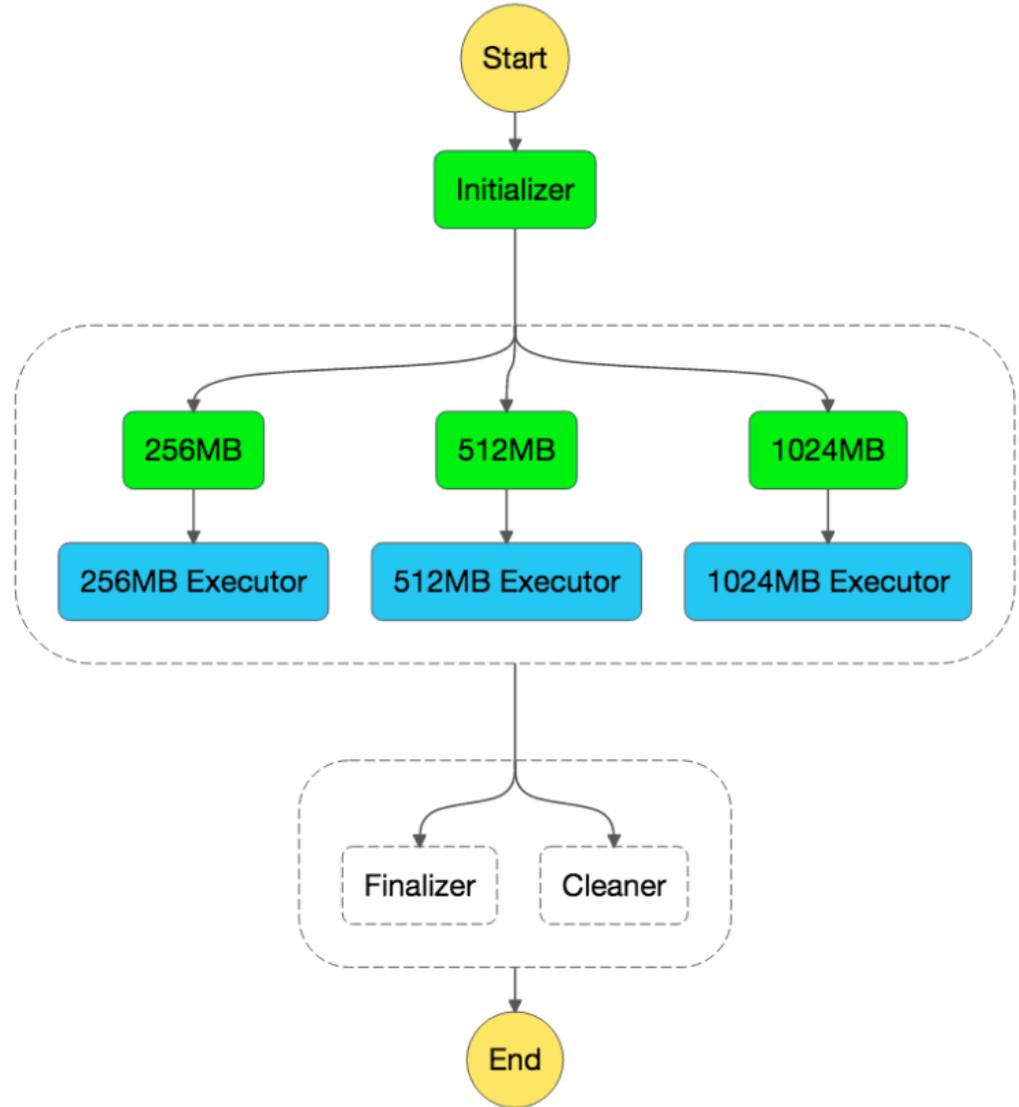
128 MB

Timeout [Info](#)

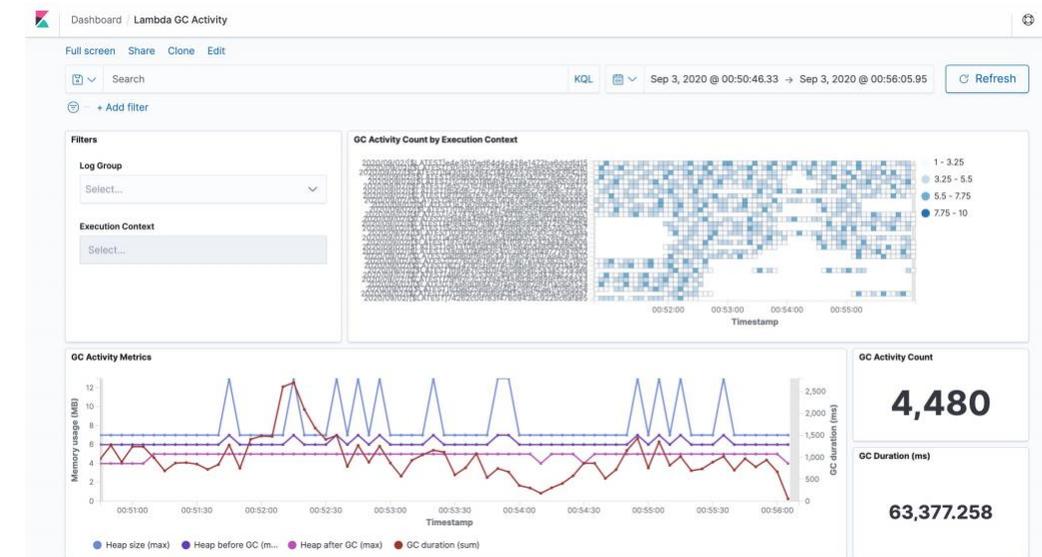
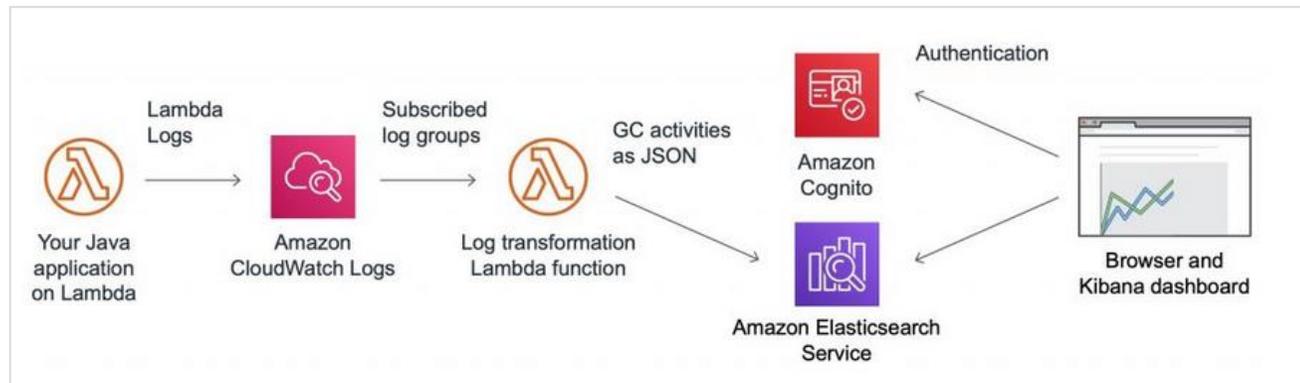
min sec

Lambda Power Tuning

- Executes different settings in parallel
- Outputs the optimal setting



Monitor the Java Virtual Machine Garbage Collection on AWS Lambda



Cost optimization

- Java is well optimized for long running server applications
 - High startup times
 - High memory utilization

And both memory and execution time are cost dimensions,
when using Serverless in the cloud

GraalVM enters the scene

GraalVM™

Project Metropolis

Goals:

Low footprint ahead-of-time mode for JVM-based languages

High performance for all languages

Convenient language interoperability and polyglot tooling

Community Edition

GraalVM Community is available for free for evaluation, development and production use. It is built from the GraalVM sources available on [GitHub](#). We provide pre-built binaries for Linux, macOS X, and Windows platforms on x86 64-bit systems. Windows support is [experimental](#).

[DOWNLOAD FROM GITHUB](#)

LICENSE

- [Open Source Licenses](#)
- Free for development and production use

BENEFITS

- Open-source license
- Free community support via [public channels](#)
- Presence of all enterprise components
- Bug fixes and enhancements

Enterprise Edition

GraalVM Enterprise provides additional performance, security, and scalability relevant for running applications in production. It is free for evaluation uses and available for download from the [Oracle Technology Network](#). We provide binaries for Linux, macOS X, and Windows platforms on x86 64-bit systems. Windows support is [experimental](#).

[DOWNLOAD FROM OTN](#)

LICENSE

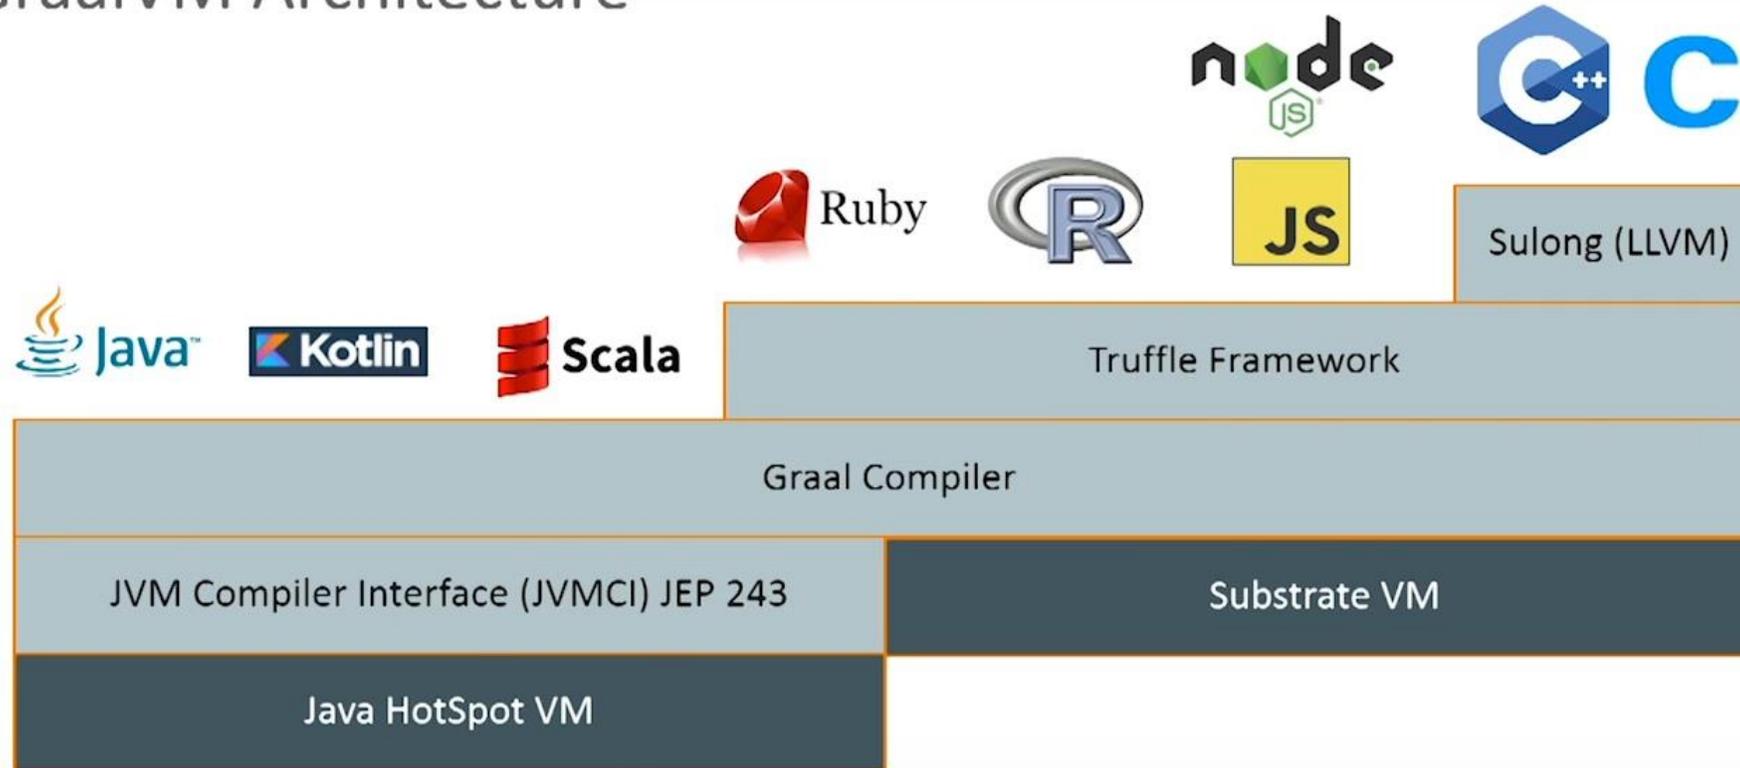
- [Oracle Master License Agreement](#)
- Free for evaluation and non-production use
- [Contact us](#) for commercial use and support options

BENEFITS

- Faster performance and smaller footprint
- Enhanced security features
- Managed capabilities for native code
- Premier 24x7x365 support via [MOS](#)

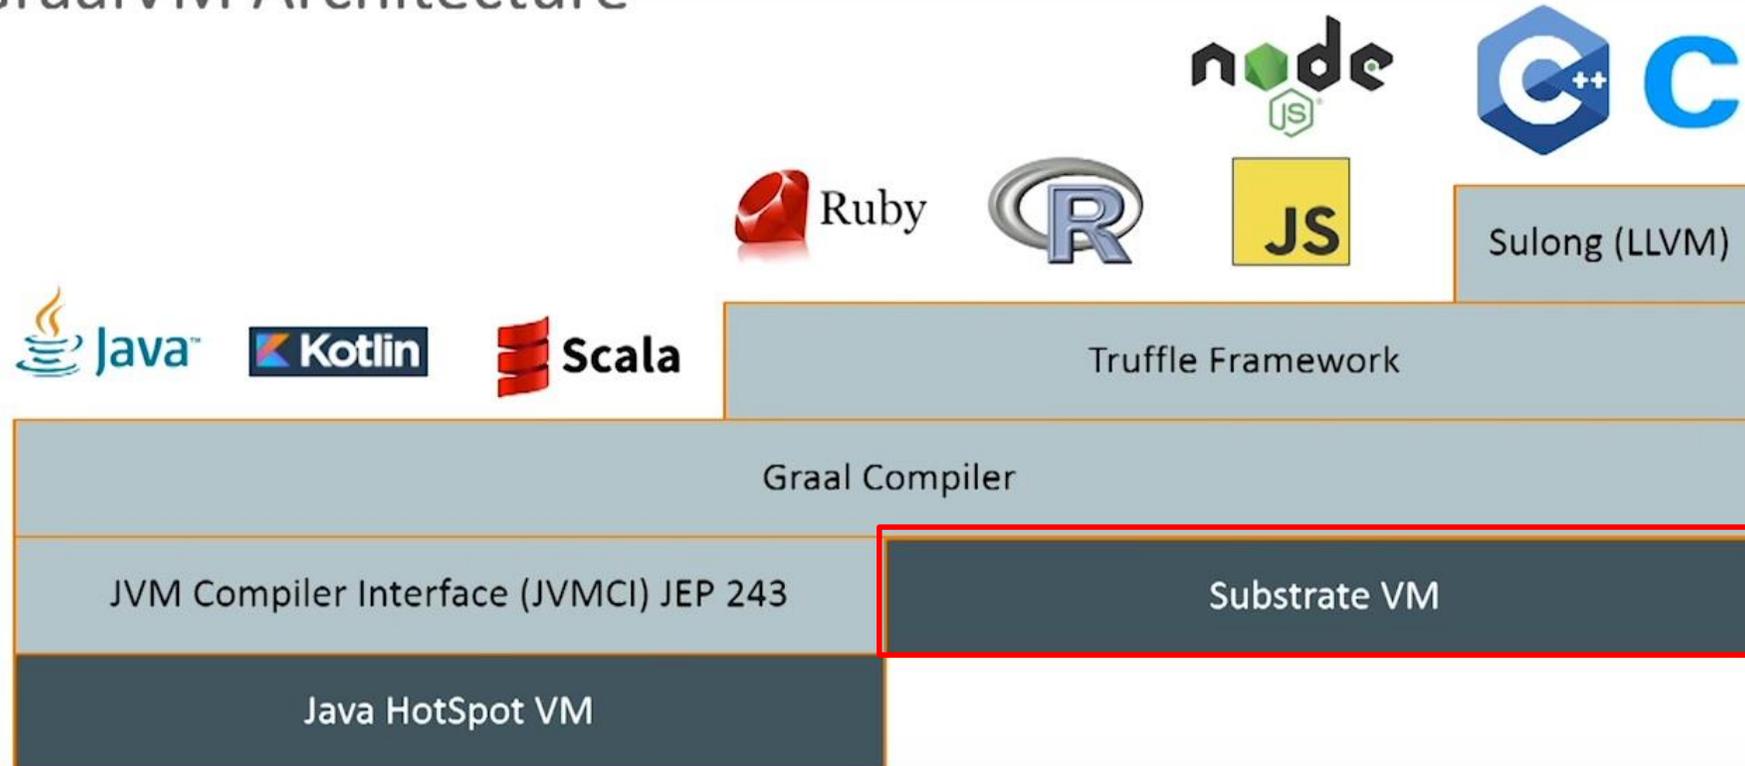
GraalVM Architecture

GraalVM Architecture

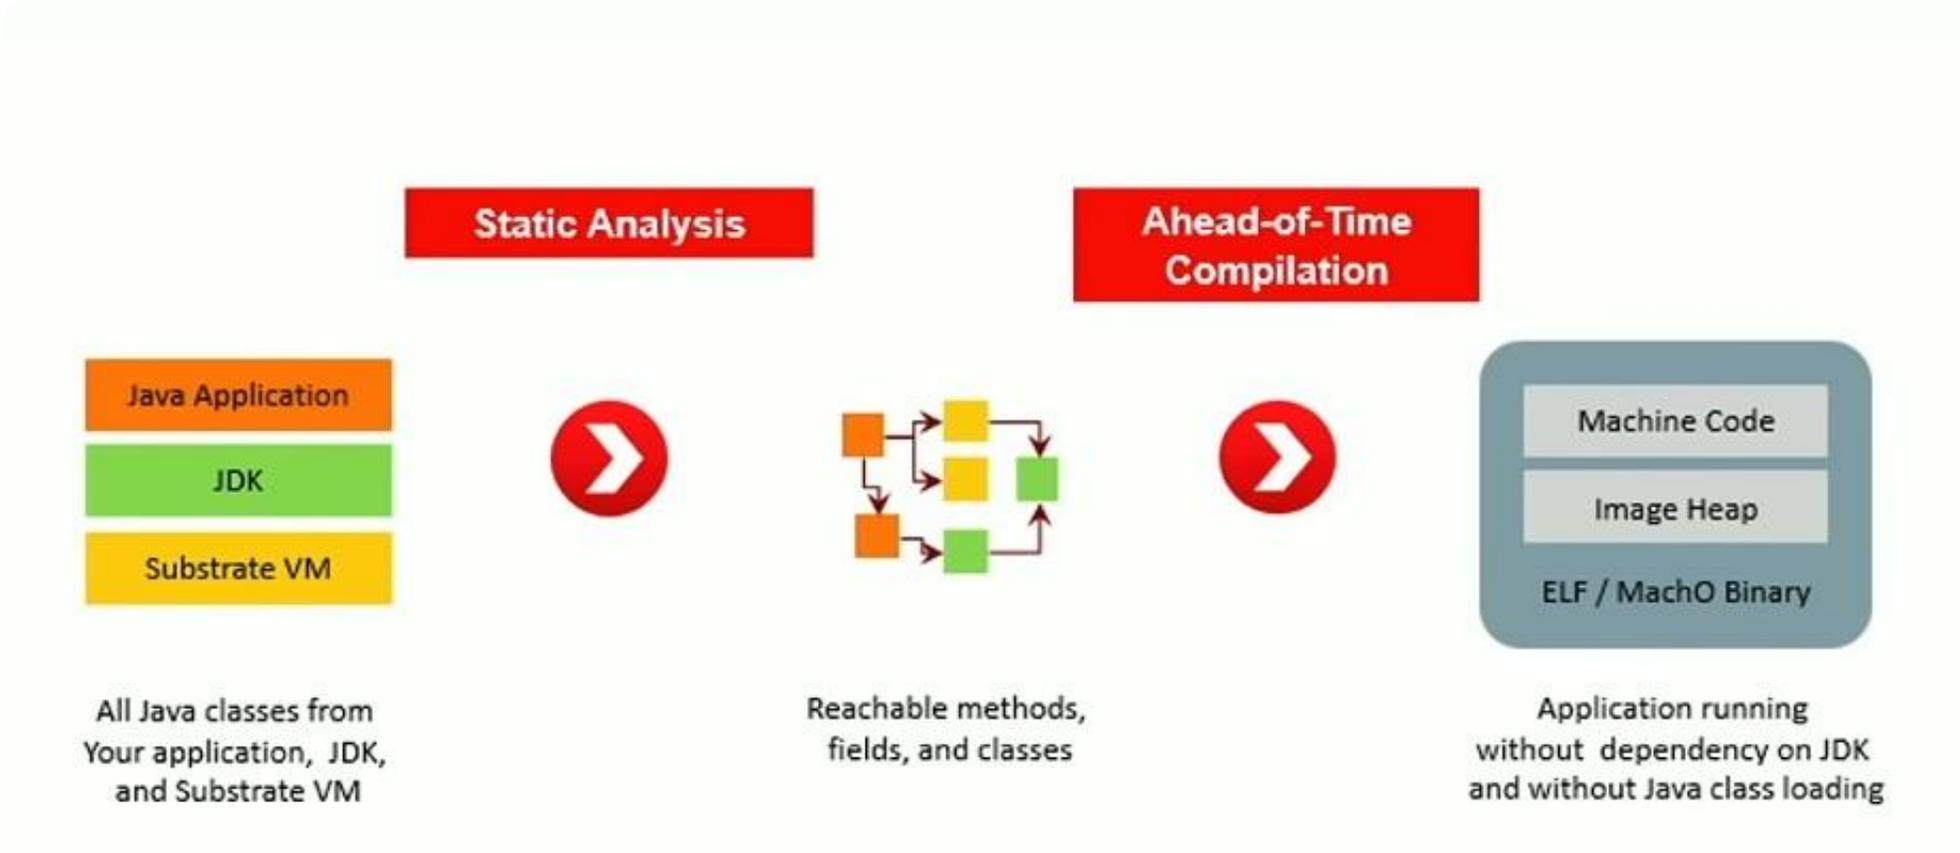


GraalVM Architecture

GraalVM Architecture



SubstrateVM

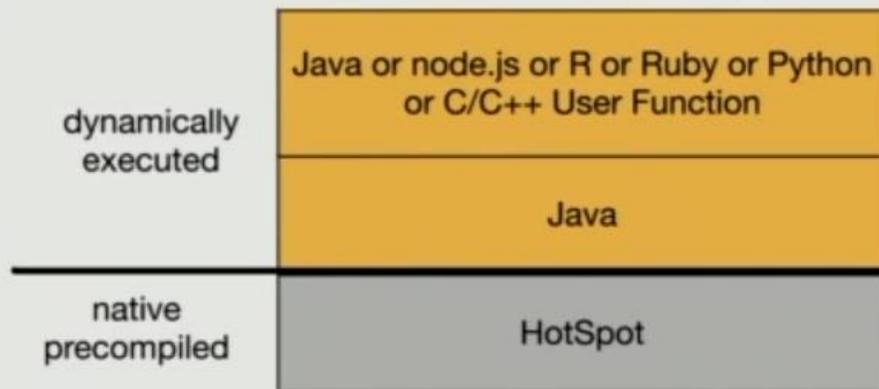


GraalVM and SubstrateVM

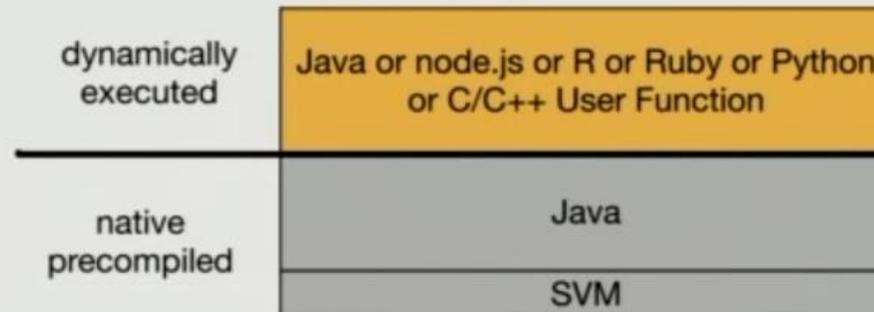
Graal VM is a hybrid of static & dynamic runtimes

- Precompile core parts of application, but still allow extensibility!

Graal VM on HotSpot



Graal VM on SubstrateVM



GraalVM on SubstrateVM

A game changer for Java & Serverless?

Java Function compiled into a **native executable** using **GraalVM on SubstrateVM** reduces

- “cold start” times
- memory footprint

by order of magnitude compared to running on JVM.

And both memory and execution time are cost dimensions, when using Serverless in the cloud

GraalVM on SubstrateVM

A game changer for Java & Serverless?

Current challenges with **native executable** using **GraalVM** :

- Most Cloud Providers (AWS) doesn't provide GraalVM as Java Runtime out of the box, only Open JDK (e.g. AWS provides Corretto)
- Some Cloud Providers (e.g. AWS) provide Custom Runtime Option

Lambda Layers & Lambda Runtime API

New for AWS Lambda – Use Any Programming Language and Share Common Components

by Danilo Poccia | on 29 NOV 2018 | in [AWS Lambda](#), [Compute](#), [Launch](#), [News](#), [Serverless](#), [Top Posts*](#) | [Permalink](#) | [Comments](#) | [Share](#)



Voiced by [Amazon Polly](#)

I remember the excitement when [AWS Lambda](#) was [announced in 2014](#)! [Four years on](#), customers are using Lambda functions for many different use cases. For example, [iRobot](#) is using AWS Lambda to [provide compute services for their Roomba robotic vacuum cleaners](#), [Fannie Mae](#) to [run Monte Carlo simulations for millions of mortgages](#), [Bustle](#) to [serve billions of requests for their digital content](#).

Today, we are introducing two new features that are going to make serverless development even easier:

- **Lambda Layers**, a way to centrally manage code and data that is shared across multiple functions.
- **Lambda Runtime API**, a simple interface to use any programming language, or a specific language version, for developing your functions.

Custom Lambda Runtimes

Custom AWS Lambda runtimes

You can implement an AWS Lambda runtime in any programming language. A runtime is a program that runs a Lambda function's handler method when the function is invoked. You can include a runtime in your function's deployment package in the form of an executable file named `bootstrap`.

A runtime is responsible for running the function's setup code, reading the handler name from an environment variable, and reading invocation events from the Lambda runtime API. The runtime passes the event data to the function handler, and posts the response from the handler back to Lambda.

Your custom runtime runs in the standard Lambda [execution environment](#). It can be a shell script, a script in a language that's included in Amazon Linux, or a binary executable file that's compiled in Amazon Linux.

To get started with custom runtimes, see [Tutorial – Publishing a custom runtime](#). You can also explore a custom runtime implemented in C++ at [awslabs/aws-lambda-cpp](#) on GitHub.

Topics

- [Using a custom runtime](#)
- [Building a custom runtime](#)

Using a custom runtime

To use a custom runtime, set your function's runtime to `provided`. The runtime can be included in your function's deployment package, or in a [layer](#).

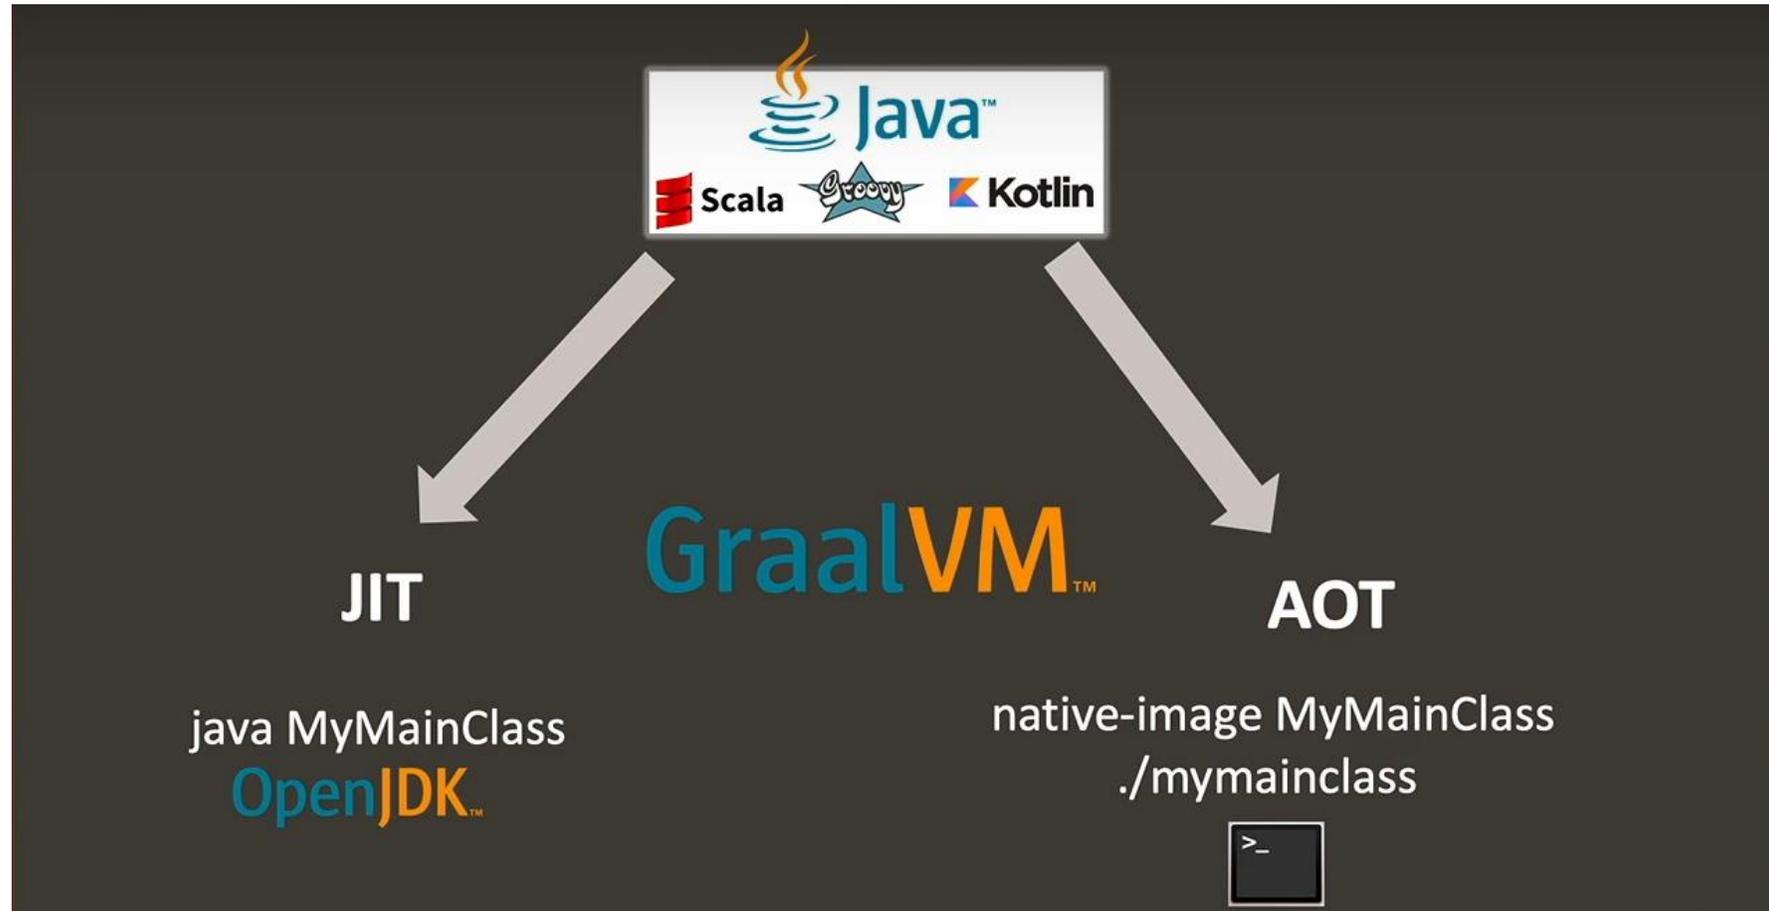
Example function.zip

```
.
├─ bootstrap
└─ function.sh
```

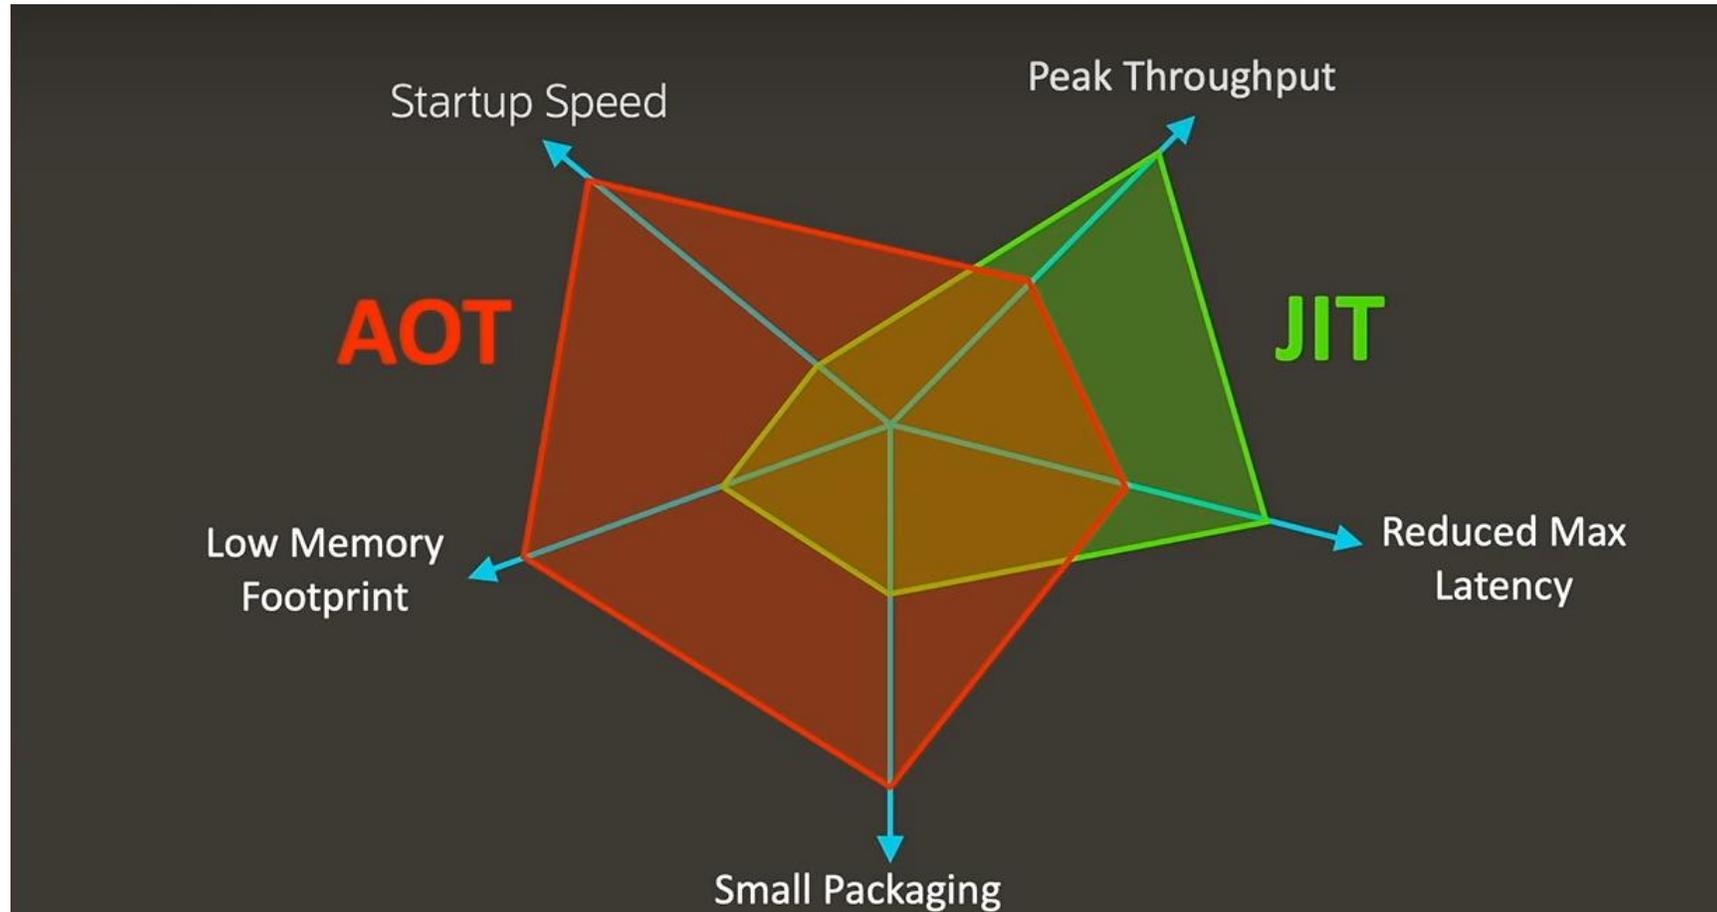


If there's a file named `bootstrap` in your deployment package, Lambda executes that file. If not, Lambda looks for a runtime in the function's layers. If the `bootstrap` file isn't found or isn't executable, your function returns an error upon invocation.

GraalVM Complitation Modes



AOT vs JIT



Source: „Everything you need to know about GraalVM by Oleg Šelajev & Thomas Wuerthinger” <https://www.youtube.com/watch?v=ANN9rxYo5Hg>

GraalVM Profile-Guided Optimizations

Profile-Guided Optimizations

GraalVM Enterprise allows to apply profile-guided optimizations (PGO) for additional performance gain and higher throughput of native images. With PGO you can collect the profiling data in advance and then feed it to the `native-image` builder, which will use this information to optimize the performance of the resulting binary.

Note: This feature is available with **GraalVM Enterprise** only.

One approach is to gather the execution profiles at one run and then use them to optimize subsequent compilation(s). In other words, you create a native image with the `--pgo-instrument` option to collect the profile information. The `--pgo-instrument` builds an instrumented native image with profile-guided optimization data collected of AOT compiled code in the `default.iprof` file, if nothing else is specified. Then you run an example program, saving the result in `default.iprof`. Finally, you create a second native image with `--pgo-profile.iprof` flag that should be significantly faster. You can collect multiple profile files and add them to the image build.

Support of GraalVM native images in Frameworks

Spring Framework: working toward **GraalVM native image support** without requiring additional configuration or workaround is one of the themes of upcoming Spring Framework 5.3

Spring Boot: Ongoing work on experimental Spring Graal Native project. Probably ready for the 2.4 release

Quarkus: a Kubernetes Native Java framework developed by Red Hat tailored for GraalVM and HotSpot, crafted from best-of-breed Java libraries and standards.

Micronaut: a modern, JVM-based, full-stack framework for building modular, easily testable microservice and serverless applications.

Steps to deploy to AWS

- Installation prerequisites
 - Framework of your choice (Micronaut, Quarkus, Spring)
 - Java 8 or 11
 - Apache Maven or Gradle
 - AWS CLI and AWS SAM CLI (for local testing)
- Build Linux executable of your application with GraalVM native-image
- Deploy Linux executable as AWS Lambda Custom Runtime
 - Function.zip with bootstrap Linux executable

Example function.zip

```
.  
├─ bootstrap  
└─ function.sh
```

AWS Lambda Deployment of Custom Runtime with SAM

```
AWS::Serverless::Function
AWSTemplateFormatVersion: '2010-09-09'
Transform: AWS::Serverless-2016-10-31
Description: AWS Serverless Micronaut API - graal.spring.demo::graal-spring-demo
Globals:
  Api:
    EndpointConfiguration: REGIONAL
Resources:
  GraalVMApiService:
    Type: AWS::Serverless::Function
    Properties:
      Handler: not.used.in.provided.runtime
      Runtime: provided
      CodeUri: native-image/function.zip
      MemorySize: 512
      Policies: AWSLambdaBasicExecutionRole
      Tracing: Active
      Timeout: 15
    Events:
      GetResource:
        Type: Api
        Properties:
          Path: /{proxy+}
          Method: any
```

Micronaut Framework



AWS Lambda with Micronaut Framework

```
import java.util.function.Function;

import javax.inject.Inject;

import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

import io.micronaut.function.FunctionBean;

@FunctionBean("monthly-invoice-generator")
public class MonthlyInvoiceGeneratorFunction
implements Function<MonthlyInvoiceRequest, MonthlyInvoiceResponse> {

    private static final Logger LOG = LoggerFactory.getLogger(MonthlyInvoiceGeneratorFunction.class);

    @Inject
    private MonthlyInvoiceGeneratorService monthlyInvoiceGeneratorService;

    @Override
    public MonthlyInvoiceResponse apply(MonthlyInvoiceRequest monthlyInvoiceRequest) {

        if (LOG.isDebugEnabled()) {
            LOG.debug("request: {}", monthlyInvoiceRequest);
        }
        return this.monthlyInvoiceGeneratorService.generateInvoice(monthlyInvoiceRequest);
    }
}
```

Testing AWS Lambda with Micronaut Framework

```
import io.micronaut.function.client.FunctionClient;□
```

```
@FunctionClient
```

```
public interface MonthlyInvoiceGeneratorClient {  
  
    @Named("monthly-invoice-generator")  
    Single<MonthlyInvoiceResponse> apply(@Body MonthlyInvoiceRequest request);  
  
}
```

```
import static org.junit.jupiter.api.Assertions.assertEquals;
```

```
import javax.inject.Inject;
```

```
import org.junit.jupiter.api.Test;
```

```
import io.micronaut.test.annotation.MicronautTest;
```

```
@MicronautTest
```

```
public class MonthlyInvoiceGeneratorFunctionTest {  
  
    @Inject  
    MonthlyInvoiceGeneratorClient client;  
  
    @Test  
    public void testMonthlyInvoiceFunction() throws Exception {  
        MonthlyInvoiceRequest monthlyInvoiceRequest = new MonthlyInvoiceRequest(9,2012, 123556);  
        assertEquals(50.5, client.apply(monthlyInvoiceRequest).blockingGet().getInvoiceAmount());  
    }  
  
}
```

Micronaut Additional Features

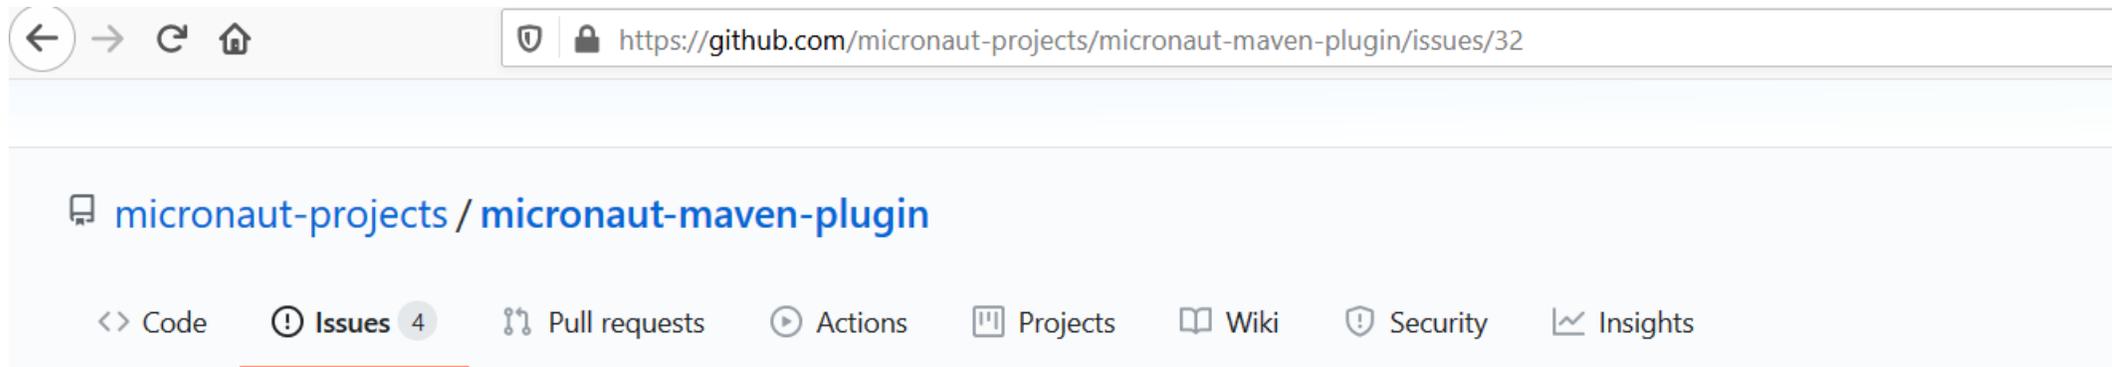
- Custom Validators
- AWS API Gateway integration
- Spring annotation processor available

```
dependencies {  
  annotationProcessor platform("io.micronaut:micronaut-bom:$micronautVersion")  
  annotationProcessor "io.micronaut:micronaut-inject-java"  
  annotationProcessor "io.micronaut:micronaut-validation"  
  annotationProcessor "io.micronaut:micronaut-graal"  
  implementation platform("io.micronaut:micronaut-bom:$micronautVersion")  
  implementation "io.micronaut:micronaut-inject"  
  implementation "io.micronaut:micronaut-validation"  
  implementation "io.micronaut:micronaut-runtime"  
  implementation "io.micronaut:micronaut-function-aws"
```

```
dependencies {  
  annotationProcessor platform("io.micronaut:micronaut-bom:$micronautVersion")  
  annotationProcessor "io.micronaut:micronaut-graal"  
  annotationProcessor "io.micronaut:micronaut-inject-java"  
  annotationProcessor "io.micronaut:micronaut-validation"  
  annotationProcessor "io.micronaut.spring:micronaut-spring-boot"  
  annotationProcessor "io.micronaut.spring:micronaut-spring-boot-annotation"  
  annotationProcessor "io.micronaut.spring:micronaut-spring-web-annotation"  
  implementation platform("io.micronaut:micronaut-bom:$micronautVersion")  
  implementation "io.micronaut:micronaut-http-client"  
  implementation "io.micronaut:micronaut-inject"  
  implementation "io.micronaut:micronaut-validation"  
  implementation "io.micronaut:micronaut-runtime"
```

Build GraalVM Native Image with Micronaut Framework

```
/usr/lib/graalvm/bin/native-image -H:+TraceClassInitialization --initialize-at-build-time=reactor.core.publisher.Mono  
--initialize-at-build-time=reactor.core.publisher.Flux --no-fallback --no-server -cp myapplication.jar
```



The screenshot shows a web browser displaying a GitHub issue page. The address bar shows the URL: <https://github.com/micronaut-projects/micronaut-maven-plugin/issues/32>. The page header includes the repository name 'micronaut-projects / micronaut-maven-plugin' and navigation links for 'Code', 'Issues 4', 'Pull requests', 'Actions', 'Projects', 'Wiki', 'Security', and 'Insights'. The 'Issues 4' link is highlighted with a red underline.

Implement GraalVM native image packaging #32

 Open alvarosanchez opened this issue 19 days ago · 0 comments

Quarkus



AWS Lambda with Quarkus Framework

```
import javax.inject.Inject;

import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

import com.amazonaws.services.lambda.runtime.Context;
import com.amazonaws.services.lambda.runtime.RequestHandler;

public class MonthlyInvoiceGeneratorFunction
implements RequestHandler<MonthlyInvoiceRequest, MonthlyInvoiceResponse> {

    private static final Logger LOG = LoggerFactory.getLogger(MonthlyInvoiceGeneratorFunction.class);

    @Inject
    private MonthlyInvoiceGeneratorService monthlyInvoiceGeneratorService;

    @Override
    public MonthlyInvoiceResponse handleRequest(MonthlyInvoiceRequest monthlyInvoiceRequest,
        final Context context) {

        if (LOG.isDebugEnabled()) {
            LOG.debug("request: {}", monthlyInvoiceRequest);
        }
        return this.monthlyInvoiceGeneratorService.generateInvoice(monthlyInvoiceRequest);
    }
}
```

Testing AWS Lambda with Quarkus Framework

```
import static org.junit.jupiter.api.Assertions.assertEquals;
import org.junit.jupiter.api.Test;

import io.quarkus.amazon.lambda.test.LambdaClient;
import io.quarkus.test.junit.QuarkusTest;

@QuarkusTest
public class MonthlyInvoiceGeneratorFunctionTest {

    @Test
    public void testMonthlyInvoiceFunction() throws Exception {
        MonthlyInvoiceRequest monthlyInvoiceRequest = new MonthlyInvoiceRequest(9,2012, 123556);
        MonthlyInvoiceResponse monthlyInvoiceResponse =
            LambdaClient.invoke(MonthlyInvoiceResponse.class, monthlyInvoiceRequest);
        assertEquals(50.5, monthlyInvoiceResponse.getInvoiceAmount());
    }
}
```

Build GraalVM Native Image with Quarkus Framework

```
<profile>  
  <id>native</id>
```

 mvn -Pnative package

```
  <activation>  
    <property>  
      <name>native</name>  
    </property>  
  </activation>  
  <build>  
    <plugins>  
      <plugin>  
        <groupId>io.quarkus</groupId>  
        <artifactId>quarkus-maven-plugin</artifactId>  
        <version>${quarkus.version}</version>  
        <executions>  
          <execution>  
            <goals>  
              <goal>native-image</goal>  
            </goals>  
            <configuration>  
              <enableHttpUrlHandler>true</enableHttpUrlHandler>  
            </configuration>  
          </execution>  
        </executions>  
      </plugin>  
    </plugins>  
  </build>  
</profile>
```

Quarkus Additional Features

- Website for creating the App
- AWS API Gateway integration
- Funqy

Quarkus Funqy

QUARKUS - FUNQY

Quarkus Funqy is part of Quarkus's serverless strategy and aims to provide a portable Java API to write functions deployable to various FaaS environments like AWS Lambda, Azure Functions, Knative, and Knative Events (Cloud Events). It is also usable as a standalone service.

Because Funqy is an abstraction that spans multiple different cloud/function providers and protocols it has to be a very simple API and thus, might not have all the features you are used to in other remoting abstractions. A nice side effect though is that Funqy is as optimized and as small as possible. This means that because Funqy sacrifices a little bit on flexibility, you'll get a framework that has little to no overhead.

Quarkus-Fanqy AWS Serverless Support

- AWS Lambda
- AWS API Gateway

Spring (Boot) Framework



Spring GraalVM Native Project

This project goal is to incubate support for building [Spring Boot](#) applications as [GraalVM native images](#).

Watch [The Path Towards Spring Boot Native Applications](#) SpringOne 2020 talk recording for more details:



It is mainly composed of:

- `spring-graalvm-native-feature` : this module is a [GraalVM feature](#). It is a kind of plugin for the native-image compilation process (which creates the native-image from the built class files). The feature participates in the compilation lifecycle, being invoked at different compilation stages to offer extra information about the application to aid in the image construction.

AWS Lambda with Spring Framework using Spring Graal Native and Spring Cloud Functions

```
import java.util.function.Function;

import org.slf4j.Logger;
import org.slf4j.LoggerFactory;
import org.springframework.beans.factory.annotation.Autowired;

public class MonthlyInvoiceGeneratorFunction
    implements Function<MonthlyInvoiceRequest, MonthlyInvoiceResponse> {

    private static final Logger LOG = LoggerFactory.getLogger(MonthlyInvoiceGeneratorFunction.class);

    @Autowired
    private MonthlyInvoiceGeneratorService monthlyInvoiceGeneratorService;

    @Override
    public MonthlyInvoiceResponse apply(MonthlyInvoiceRequest monthlyInvoiceRequest) {

        if (LOG.isDebugEnabled()) {
            LOG.debug("request: {}", monthlyInvoiceRequest);
        }
        return this.monthlyInvoiceGeneratorService.generateInvoice(monthlyInvoiceRequest);
    }
}
```

Bean Registration with Spring Framework using Spring Graal Native and Spring Cloud Functions

```
import org.springframework.boot.autoconfigure.SpringBootApplication;
import org.springframework.cloud.function.context.FunctionRegistration;
import org.springframework.cloud.function.context.FunctionType;
import org.springframework.cloud.function.context.FunctionalSpringApplication;
import org.springframework.context.ApplicationContextInitializer;
import org.springframework.context.support.GenericApplicationContext;

@SpringBootApplication(proxyBeanMethods = false)
public class DemoApplication implements ApplicationContextInitializer<GenericApplicationContext> {

    public static void main(String[] args) {
        FunctionalSpringApplication.run(DemoApplication.class, args);
    }

    @Override
    public void initialize(GenericApplicationContext context) {

        context.registerBean("monthlyInvoiceGenerator", FunctionRegistration.class,
            () -> new FunctionRegistration<>(new MonthlyInvoiceGeneratorFunction()).
                type(FunctionType.from(MonthlyInvoiceRequest.class).
                    to(MonthlyInvoiceResponse.class)));
    }
}
```

Build GraalVM Native Image with Spring Framework

mvn **-Pnative** package

```
<profile>
  <id>native</id>
  <build>
    <plugins>
      <plugin>
        <groupId>org.graalvm.nativeimage</groupId>
        <artifactId>native-image-maven-plugin</artifactId>
        <version>20.2.0</version>
        <configuration>
          <mainClass>com.example.demo.DemoApplication</mainClass>
          <buildArgs>-J-Xmx4G -H:+TraceClassInitialization -H:+ReportExceptionStackTraces
            -Dspring.native.remove-unused-autoconfig=true -Dspring.native.remove-yaml-support=true
          </buildArgs>
          <additionalBuildArgs>--delay-class-initialization-to-runtime=sun.nio.ch.WindowsAsynchronousFileChannelImpl
            </additionalBuildArgs>
          <imageName>${project.artifactId}</imageName>
        </configuration>
        <executions>
          <execution>
            <goals>
              <goal>native-image</goal>
            </goals>
            <phase>package</phase>
          </execution>
        </executions>
      </plugin>
    </plugins>
  </build>
</profile>
```

Framework Comparison

- Project Initializer
- Programming Model
- Database Support
- Test Support
- Native Image
 - Native image size
 - Startup time
 - Heap size

Conclusion

- GraalVM and Frameworks are really powerful with a lot of potential
- But in combination with Native Image currently not without challenges
 - AWS Lambda Custom runtime requires Linux executable only
 - Windows and Mac developers may only build Linux executable via Docker
 - plenty of *'No instances of ... are allowed in the image heap'* and other errors when building a native image
 - Lots of experimentation with additional build arguments like “initialize-at-runtime” or “delay-class-initialization-to-runtime“ required
- Once again: **Less (dependencies, classes) is more**
 - AWS Lambda function should be small and shouldn't have many dependencies

Try it yourselves

- Micronaut
 - <https://github.com/micronaut-guides/micronaut-function-aws-lambda>
- Quarkus
 - <https://github.com/JosemyDuarte/quarkus-terraform-lambda-demo/tree/dynamo-terraform>
- Spring Boot
 - <https://github.com/spring-projects-experimental/spring-graalvm-native/tree/master/spring-graalvm-native-samples/function-aws>
- Misc
 - <https://github.com/awslabs/aws-serverless-java-container/tree/master/samples>

Project Leyden

Call for Discussion: New Project: Leyden

mark.reinhold at oracle.com [mark.reinhold at oracle.com](mailto:mark.reinhold@oracle.com)

Mon Apr 27 16:38:55 UTC 2020

- Previous message: [Type-parameterized complement to Object.equals\(Object\)](#)
- Next message: [Call for Discussion: New Project: Leyden](#)
- Messages sorted by: [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

I hereby invite discussion of a new Project, Leyden, whose primary goal will be to address the long-term pain points of Java's slow startup time, slow time to peak performance, and large footprint.

Leyden will address these pain points by introducing a concept of `_static images_` to the Java Platform, and to the JDK.

- A static image is a standalone program, derived from an application, which runs that application -- and no other.
- A static image is a closed world: It cannot load classes from outside the image, nor can it spin new bytecodes at run time.

These two constraints enable build-time analyses that can remove unused classes and identify class initializers which can be run at build time, thereby reducing both the size of the image and its startup time. These constraints also enable aggressive ahead-of-time compilation, thereby reducing the image's time to peak performance.





Thank You!