


Java 2022: Containerized, Serverless, Cloud-native

01


\$ whoami

- developer    ...
- devops guy 
- trainer
- speaker   ...


Timeline



Wall of confusion




Devs




Ops


Wall of confusion




No-walls/Shift-left



Wall of new confusion




Devs



DevOps

Wall of new confusion



Definitions



Cloud-native

“

*Cloud native technologies empower organizations to build and run **scalable** applications in **modern, dynamic environments** such as public, private, and hybrid clouds. Containers, service meshes, micro-services, immutable infrastructure, and declarative APIs exemplify this approach.*



Cloud-native

“

*These techniques enable **loosely coupled systems** that are **resilient, manageable, and observable**. Combined with robust automation, they allow engineers to make **high-impact changes frequently and predictably with minimal toil**.*



Serverless

“


I don't want to care about servers.

cloud-native ≠ running in the cloud

serverless ≠ no servers

containerized ≠ running in docker

12-factor application



14


IV. backing services =
network resources

API-first


Telemetry

Authn and authz

Secret management



vertical




cpu +4
mem +16


horizontal




machines +2

redesign









J2EE

JEE

JakartaEE

MicroProfile

monolithic release
proprietary control


composable
community-driven
independently evolving



multi-core
many GBs



1 vCPU
64 MB



crate = VM | container | function

**Make JAR, not
WAR**

Fat-less app




Container-awareness

“

Inside Linux containers, OpenJDK versions 8 and later can correctly detect the container-limited number of CPU cores and available RAM. For all currently supported OpenJDK versions this is turned on by default.

memory areas



J7≤



J8



J11



J17



29

CPU cores

RAM memory

container

JVM

GC threads, JIT threads,
various framework threads,
memory pools sizes etc.



Container support

- Container support is enabled by default since Java 10
- Some settings and parameters are back-ported to Java 8
- Can be disabled with `-XX:-UseContainerSupport`




Container support

- `-XX:ActiveProcessorCount=count`
- `-XX:InitialRAMPercentage=mem`
- `-XX:MaxRAM=mem`
- `-XX:MaxRAMPercentage=pct`
- `-XX:MinRAMPercentage=pct`

GC


1791

Image size



JDK


JPMS



~290MB (uncompressed)


≥J9

JDK Tangle



java.base	jdk.compiler	jdk.jlink
java.compiler	jdk.crypto.cryptoki	jdk.jpackage
java.datatransfer	jdk.crypto.ec	jdk.jshell
java.desktop	jdk.crypto.mscapi	jdk.jsobject
java.instrument	jdk.dynalink	jdk.jstard
java.logging	jdk.editpad	jdk.localedata
java.management	jdk.hotspot.agent	jdk.management
java.management.rmi	jdk.httpserver	jdk.management.agent
java.naming	jdk.incubator.foreign	jdk.management.jfr
java.net.http	jdk.incubator.vector	jdk.naming.dns
java.prefs	jdk.internal.ed	jdk.naming.rmi
java.rmi	jdk.internal.jvmstat	jdk.net
java.scripting	jdk.internal.le	jdk.nio.mapmode
java.se	jdk.internal.opt	jdk.random
java.security.jgss	jdk.internal.vm.ci	jdk.sctp
java.security.sasl	jdk.internal.vm.compiler	jdk.security.auth
java.smartcardio	jdk.internal.vm.compiler.management	jdk.security.jgss
java.sql	jdk.jartool	jdk.unsupported
java.sql.rowset	jdk.javadoc	jdk.unsupported.desktop
java.transaction.xa	jdk.jcmd	jdk.xml.dom
java.xml	jdk.jconsole	jdk.zipfs
java.xml.crypto	jdk.jdeps	
jdk.accessibility	jdk.jdi	
jdk.attach	jdk.jdwp.agent	
jdk.charsets	jdk.jfr	

~70




~290MB (uncompressed)

~170MB (compressed)


39

~39MB (uncompressed)

~13MB (compressed)



attack surface



Log4shell




The log4j JNDI Attack

and how to prevent it

An attacker inserts the JNDI lookup in a header field that is likely to be logged.

The string is passed to log4j for logging

log4j interpolates the string and queries the malicious LDAP server.





Observability

“

In software, observability is the ability to ask new questions of the health of your running services without deploying new instrumentation.




The three pillars

- Metrics
- Logs
- Traces

Prometheus

Prometheus architecture



Exporters

JMX exporter



JMX exporter

https://github.com/prometheus/jmx_exporter



JMX exporter

```
01. java -javaagent:./jmx_prometheus_javaagent-0.16.1.jar=8080:config.
```

Spring Actuator



[◀ Back to index](#)

- 1. Enabling Production-ready Features
- 2. Endpoints
- 3. Monitoring and Management over HTTP
- 4. Monitoring and Management over JMX
- 5. Loggers
- 6. Metrics
 - 6.1. Getting started
 - 6.2. Supported Monitoring Systems
 - 6.2.1. AppOptics
 - 6.2.2. Atlas
 - 6.2.3. Datadog
 - 6.2.4. Dynatrace
 - 6.2.5. Elastic
 - 6.2.6. Ganglia
 - 6.2.7. Graphite
 - 6.2.8. Humio

By default, metrics are published via REST calls but it is also possible to use the Java Agent API if you have it on the classpath:

A screenshot of a configuration interface showing a "Properties" tab selected. It displays the line: "management.metrics.export.newrelic.client-provider-type=insights-agent". A "PROPERTIES" link is visible in the top right corner of the panel.

Finally, you can take full control by defining your own `NewRelicClientProvider` bean.

6.2.13. Prometheus

Prometheus expects to scrape or poll individual app instances for metrics. Spring Boot provides an actuator endpoint available at `/actuator/prometheus` to present a Prometheus `scrape` with the appropriate format.



The endpoint is not available by default and must be exposed, see [exposing endpoints](#) for more details.

Here is an example `scrape_config` to add to `prometheus.yml`:

Akka



Lightbend
DOCUMENTATION

- ▶ Introduction
- ▶ Getting started
- ▶ Developer sandbox
- ▶ Setup
- ▶ Instrumentations
- ▶ Extensions
- ▼ **Backend plugins**
 - Cinnamon metadata
 - Datadog
 - New Relic
 - Coda Hale Metrics
- ▼ **Prometheus**
 - Cinnamon dependency
 - Exporters
 - Configuration
 - Custom exporter

Docs / Lightbend Telemetry / Backend plugins / Prometheus

LANGUAGES **Scala** ▾ BUILD TOOLS **Maven** ▾

Prometheus

Lightbend Telemetry can report metrics to [Prometheus](#), using a backend plugin integrated with the [Prometheus JVM client](#).

Cinnamon dependency

First make sure that your build is configured to use the [Cinnamon Agent](#) and has instrumentations enabled, such as [Akka instrumentation](#) or [Akka HTTP instrumentation](#).

Here is the core Cinnamon Prometheus dependency, but note that you also need to select an [exporter](#).



Prometheus Client for Java

```
01. import io.prometheus.client.Counter;  
  
class YourClass {  
  
    static final Counter requests = Counter.build()  
        .name("requests_total").help("Total requests.").register();  
  
    void processRequest() {  
  
        requests.inc();  
  
        // Your code here.  
  
    }  
}
```



Traces



Definition


A trace is a data/execution path through the system, and can be thought of as a directed acyclic graph of spans.

Trace/Span




Jaeger

Jaeger architecture



Instrumentation





Get tracer

```
01. Tracer tracer = Configuration.fromEnv().getTracer();
```



Access span

```
01. Span span = tracer.scopeManager().activeSpan();  
02. if (span != null) {  
03.     span.log("...");  
04. }
```

Create span

```
01. Span span = tracer.buildSpan("someWork").start();  
02. try (Scope scope = tracer.scopeManager().activate(span)) {  
03.     // Do things.  
04. } catch(Exception ex) {  
05.     Tags.ERROR.set(span, true);  
06.     span.log(Map.of(Fields.EVENT, "error", Fields.ERROR_OBJECT,  
07.                     ex, Fields.MESSAGE, ex.getMessage()));  
08. } finally {  
09.     61     span.finish();  
10. }
```



Ignore parent span

```
01. Span span = tracer.  
02.           buildSpan("someWork").  
03.           ignoreActiveSpan().  
04.           start();
```

JFR

Java Flight Recorder

- It was first introduced in JRockit.
- Many features of JRockit including JFR were merged into Oracle HotSpot at version 8.
- Till version 11, JFR/JMC was considered a commercial feature (-
`XX:+UnlockCommercialFeatures -XX:+FlightRecorder`).
- In 11, JFR became free, but JMC (Mission Control UI) was removed from JDK, but remained a separate utility.



jcmd

01. jcmb <PID> JFR.start duration=60s filename=recording.jfr
02. jcmb <PID> JFR.start
03. jcmb <PID> JFR.dump name=1 filename=recording.jfr
04. jcmb <PID> JFR.stop



jfr

01. jfr print --events CPULoad,GarbageCollection recording.jfr
02. jfr print --categories "GC,JVM,Java*" recording.jfr
03. jfr summary recording.jfr
04. jfr metadata recording.jfr



Custom events

```
01. import jdk.jfr.Event;  
02. public class RestCallEvent extends Event {  
03.     public String path;  
04.     public String key;  
05.     public long dataSize;  
06.  
07. }
```



Custom events

```
01. event.begin();
02. // do something
03. event.key = key;
04. event.dataSize = val.length();
05. // do something
06. event.end();
07. event.commit();
```

JFR Streaming (Java 14)

```
01. try (var rs = new RecordingStream()) {  
02.     rs.enable("jdk.CPULoad").withPeriod(Duration.ofSeconds(1));  
03.     rs.enable("jdk.JavaMonitorEnter").withThreshold(Duration.ofMilli  
04.     rs.onEvent("jdk.CPULoad", event -> {  
05.         System.out.println(event.getFloat("machineTotal"));  
06.     });  
07.     rs.onEvent("jdk.JavaMonitorEnter", event -> {  
08.         System.out.println(event.getClass("monitorClass"));  
09.     }69);  
10.    rs.start();
```

JFR Streaming (Java 14)

```
01. Configuration c = Configuration.getConfiguration("default");  
02. try (var rs = new RecordingStream(c)) {  
03.     rs.onEvent("jdk.GarbageCollection", System.out::println);  
04.     rs.onEvent("jdk.CPULoad", System.out::println);  
05.     rs.onEvent("jdk.JVMInformation", System.out::println);  
06.     rs.start();  
07. }  
08. }
```

AWS CodeGuru

AWS X-ray

Serverless



If it were 2005...

- Jetty/Tomcat/Ruby/PHP (on a server)
- MySQL, PostgreSQL, HSQL, Sqlite (on a server)
- File system (on a server)



If it were 2015...

- Jetty/Tomcat/Ruby/PHP (in a container on a server) or PaaS in the cloud
- MySQL, PostgreSQL, HSQL, Sqlite (on a server or in a container on a server) or DaaS in the cloud
- File system (in a volume on a server) or object storage in the cloud



In 2022...

- CDN + FaaS + LB in the cloud
- DaaS in the cloud
- Object storage in the cloud

FaaS JVM choices

GCP



GCP (512m)

01. -XX:MaxRAM=512m
02. -XX:MaxRAMPercentage=70



GCP (512m)

- Copy
- MarkSweepCompact



GCP (4096m)

01. -XX:MaxRAM=4096m
02. -XX:MaxRAMPercentage=70



GCP (4096m)

- G1 Young
- G1 Old

Azure



Azure

01. -XX:+TieredCompilation
02. -XX:TieredStopAtLevel=1
03. -Xverify:none
04. -Djava.net.preferIPv4Stack=true



GC

- PS Scavenge
- PS MarkSweep



AWS (512)

01. -XX:MaxHeapSize=445645k
02. -XX:MaxMetaspaceSize=52429k
03. -XX:ReservedCodeCacheSize=26214k
04. -XX:+UseSerialGC



AWS (4096)

01. -XX:MaxHeapSize=3948544k
02. -XX:MaxMetaspaceSize=163840k
03. -XX:ReservedCodeCacheSize=81920k
04. -XX:+UseSerialGC

AWS

- 
01. -javaagent:/var/runtime/amzn-log4j-security-jdk11-0.1alpha.jar
 02. -Xshare:on
 03. -XX:SharedArchiveFile=/var/lang/lib/server/runtime.jsa
 04. -XX:-TieredCompilation
 05. -Djava.net.preferIPv4Stack=true



CDS

- Class Data Sharing
- It contains 1300+ core library classes loaded by the bootstrap class loader
- It is stored in a format that can be loaded very quickly, compared to loading from a JAR file



Dynamic CDS

01. -XX:ArchiveClassesAtExit=cds.jsa



JVM FaaS on clouds

- AWS: Serial, only JIT C2, shared class data, memory set explicitly
- GCP: Serial or G1, auto-tuning memory pools
- Azure: Parallel, only JIT C1

JIT

JVMCI



Graal Compiler

- Graal Compiler = JIT Compiler written in Java
- Added in Java 10
- Removed in Java 17



Graal VM

- OpenJDK with Graal compiler
- Truffle framework
- Tooling for other languages (Python, Node.js, Ruby etc.)
- Native image + Substrate VM



Modes

- JVM
- Native

Native



Nice and shiny?

- Requires extra tooling (platform image, C++ compiler)
- Compilation time could be quite long for larger code bases
- "Closed-world" approach requires extra configuration for handling reflection, dynamic proxies etc.
- Some runtime tooling is not available (no way to get a memory dump)



Reflection

01. native-image -H:ReflectionConfigurationFiles=r.json ...



Reflection

```
01. {  
02.   {  
03.     "name":  
04.       "java.lang.String$CaseInsensitiveComparator",  
05.     "queriedMethods": [  
06.       { "name": "compare" }  
07.     ]  
08.   }  
09. }
```

100



Hmmm...

- Long compilation times
- Extra configuration or code changes are required
- Not all Java functionality is supported



Why bother?

- improved startup time (= faster scaling)
- reduced memory usage (= cheaper)
- reduced image size
- better performance (?)
- better security (?)
- ideal for serverless/ML workloads



Community!

- **Spring Native** (native image support for Spring/Spring Boot)
- **Quarkus** (RedHat's baby, lots of integrations, community work, build-time code generation)
- **Micronaut** (no reflection, build-time code generation)
- **Helidon** (Oracle's baby, support for jlink)

Quarkus

Helidon

Infra-as-code

CDK



Infra-as-Java

```
01. Bucket bucket = Bucket.Builder  
02.     .create(this, targetBucket).build();  
03. PolicyStatement statement1 = PolicyStatement.Builder.create()  
04.     .effect(Effect.ALLOW)  
05.     .actions(asList("s3:GetBucket", "s3:PutObject"))  
06.     .resources(asList("arn:aws:s3:::" + bucket.getBucketName() + "/*")  
07.     .build();  
08.     .build();
```



Infra-as-Java

```
01. Vpc vpc = new Vpc(this, "VPC");  
02. AutoScalingGroup asg = AutoScalingGroup.Builder  
04.     .create(this,"ASG")  
05.     .vpc(vpc)  
06.     .instanceType(InstanceType.of(BURSTABLE2, MICRO))  
07.     .machineImage(new AmazonLinuxImage())  
08.     .build();
```



Infra-as-Java


```
01. HealthCheck.Builder healthCheckBuilder =  
02.     new HealthCheck.Builder();  
03. HealthCheck healthCheck = healthCheckBuilder.port(80).build();  
04. LoadBalancer lb = LoadBalancer.Builder.create(this,"LB")  
05.     .vpc(vpc)  
06.     .internetFacing(Boolean.TRUE)  
07.     .healthCheck(healthCheck)  
08.     .build();
```



Conclusion

- JVM ecosystem has many options for different load types.
- It makes it useful for any cloud-native/serverless/containerized environment.
- Community is vibrant and responsive.

Golden Path



Future



Project Loom

```
01. Thread.startVirtualThread(  
02.     () -> {  
03.         System.out.println("Hello World");  
04.     }  
05. );
```



Alibaba's Dragonwell

<https://dragonwell-jdk.io/>



Project Leyden

“

The primary goal of this Project is to address the long-term pain points of Java's slow startup time, slow time to peak performance, and large footprint.

Thank you!

Questions?

118

\$ ping me

 @codingandrey

 github.com/aadamovich

 lv.linkedin.com/in/andreyadamovich

 extremeautomation.io

 andrey@extremeautomation.io