



Introduction to JVM Architecture

Learn about the architecture of the Java Virtual Machine (JVM)

[Get started](#)

Overview

This course provides an introduction to the architecture of the Java Virtual Machine (JVM). You will learn about the various components of the JVM and how they work together to execute Java programs. By understanding the JVM architecture, you will be able to write more efficient and optimized Java code.

Overview of JVM Architecture

01 | Overview of JVM Architecture

Overview of JVM Architecture

Introduction

The Java Virtual Machine (JVM) is an integral part of the Java platform, providing an execution environment for Java applications. It is responsible for running compiled Java bytecode, translating it into machine-specific instructions that can be executed by the underlying system. Understanding the architecture of the JVM is essential for Java developers as it enables them to write efficient and portable code. In this topic, we will delve into the details of the JVM architecture.

Class Loader Subsystem

One of the key components of the JVM architecture is the Class Loader subsystem. It is responsible for loading classes and interfaces during the runtime

of a Java program. The Class Loader subsystem consists of three distinct components:

1. **Bootstrap Class Loader:** This is the first component in the Class Loader subsystem and is responsible for loading core Java classes from the bootstrap classpath. These classes are crucial for initializing the JVM.
2. **Extension Class Loader:** The extension class loader is responsible for loading classes from the extension classpath. It loads classes that provide additional functionality to the core Java platform.
3. **Application Class Loader:** Also known as the system class loader, it loads classes from the application classpath, which includes user-defined classes and libraries.

Runtime Data Areas

The JVM architecture includes several runtime data areas that are used during the execution of a Java program. These data areas store the necessary information for the JVM to execute bytecode efficiently. The main runtime data areas are:

1. **Method Area:** The method area stores class-level data shared across all threads, including the bytecode itself, constant pool, field and method data, and runtime constant pool.
2. **Heap Area:** The heap area is where objects are allocated at runtime. It is divided into multiple generations (young, old, and permanent) to optimize garbage collection.
3. **Stack Area:** Each thread of execution in a Java program has its own stack, which stores method-specific data, including local variables and method invocations. The stack is organized in frames that correspond to individual method calls.
4. **PC Registers:** The program counter (PC) registers store the address of the current instruction being executed by a thread. Each thread has its own PC register.
5. **Native Method Stacks:** The native method stack is used for executing native (non-Java) code. It is distinct from the Java stack and contains information about native methods.

Execution Engine

The execution engine is responsible for executing the instructions of Java bytecode. It consists of two components:

1. **Interpreter:** The interpreter reads and interprets bytecode instructions one by one, executing them directly. While this approach is straightforward, it can be slower compared to other execution modes.
2. **Just-In-Time (JIT) Compiler:** The JIT compiler dynamically translates bytecode into native machine code, optimizing it for the underlying hardware. This approach improves execution speed by reducing interpretation overhead.

Garbage Collection

Memory management is a crucial aspect of any programming language, and the JVM architecture includes a garbage collector for automatic memory management. The garbage collector, or GC, monitors the heap area and identifies objects that are no longer in use. It reclaims the memory occupied by these objects, freeing it up for future allocation.

Conclusion

In this topic, we explored the architecture of the JVM, focusing on its components and runtime data areas. We discussed the Class Loader subsystem, runtime data areas such as the method area, stack area, and heap area, the execution engine, and the garbage collector. Understanding the JVM architecture provides valuable insights into how Java programs are executed, enabling developers to write efficient and portable code.

Class Loading and Memory Management in JVM

02 | Class Loading and Memory Management in JVM

Class Loading in JVM

Class loading is the process by which Java classes are loaded and initialized within the Java Virtual Machine (JVM). When a Java program is executed, the JVM loads the required classes into memory before they can be executed. Class loading involves three main steps:

- 1. Loading:** During this phase, the JVM searches for the binary representation (bytecode) of a class and reads it into memory. The bytecode can be obtained from a local file system, network, or any other source. At this stage, the JVM creates a unique runtime representation of the class called a `Class` object.
- 2. Linking:** This phase involves three sub-stages: verification, preparation, and resolution. During verification, the JVM ensures that the loaded bytecode is valid and does not violate any security constraints. The preparation phase allocates memory for static variables and initializes them with default values. Resolution resolves symbolic references to other classes or methods.
- 3. Initialization:** This is the final phase where the JVM executes the static initializers (also known as static blocks) in the class. Static initializers are used to initialize static variables

and perform any additional setup required by the class. The initialization process is performed automatically when the class is first used.

The class loading process is dynamic, allowing classes to be loaded and unloaded during runtime, providing flexibility and dynamic behavior to Java programs.

Memory Management in JVM

Memory management in the JVM refers to the allocation, deallocation, and usage of memory during the execution of a Java program. JVM's memory is divided into different regions, namely the method area, heap, and stack.

1. **Method Area:** The method area stores class-level data, including the bytecode of loaded classes, method code, constant pool, and static variables. Each loaded class has its own runtime constant pool, which is a table of symbolic references used internally by the runtime.
2. **Heap:** The heap is the runtime data area where objects and their instance variables are allocated. Memory for objects is dynamically allocated on the heap, and the JVM automatically reclaims memory from objects that are no longer referenced, through a process known as garbage collection.
3. **Stack:** The stack is used for storing method-specific data and local variables. Each thread in the JVM has its own stack, which stores temporary data during method invocation and execution. The stack also keeps track of method calls, allowing for method entry and exit.

The garbage collector, a critical component of memory management, automatically identifies and reclaims memory from objects that are no longer reachable, freeing up memory for future allocations. It follows different garbage collection algorithms, such as mark-and-sweep, moving, or generational collection, to efficiently manage memory.

Proper memory management ensures efficient utilization of resources, prevents memory leaks, and enhances the overall performance of Java applications.

Execution Engine and Just-In-Time (JIT) Compilation in JVM

03 | Execution Engine and Just-In-Time (JIT) Compilation in JVM

Introduction

The Java Virtual Machine (JVM) is a crucial component in executing Java programs. It provides a platform-independent environment that allows Java code to run on any operating system. Understanding the architecture of the JVM is essential for Java developers to optimize their code and improve performance. In this topic, we will delve into the Execution Engine and Just-In-Time (JIT) Compilation in the JVM.

Execution Engine

The Execution Engine is the component responsible for executing bytecode instructions generated by the Java compiler. It interprets these instructions and produces the corresponding output. The Execution Engine consists of three main

sub-components: the Class Loader, the Runtime Data Areas, and the Execution System.

Class Loader

The Class Loader is responsible for loading Java classes into the JVM. It locates and reads the necessary class files from the file system or network, and then creates the corresponding Java class representations in memory. There are three types of Class Loaders in the JVM: Bootstrap Class Loader, Extensions Class Loader, and Application Class Loader. Each Class Loader has its own class loading hierarchy.

Runtime Data Areas

The Runtime Data Areas are the memory areas used by the JVM at runtime. These areas include the Method Area, Heap, Java Virtual Machine Stacks, and the PC Registers.

- The Method Area stores class-level data, such as the runtime constant pool, field and method data, and method code.
- The Heap is where objects are dynamically allocated.
- The Java Virtual Machine Stacks contain method frames that store local variables and partial results.
- The PC Registers store the address of the currently executing instruction.

Execution System

The Execution System is responsible for executing the bytecode instructions. There are two execution modes: interpretation and Just-In-Time (JIT) Compilation.

Just-In-Time (JIT) Compilation

Just-In-Time (JIT) Compilation is a technique used by the JVM to improve the performance of Java programs. In the JIT compilation process, the bytecode instructions are dynamically compiled into native machine code before being executed by the CPU. This allows the JVM to optimize the performance by identifying and optimizing frequently executed code segments.

JIT Compilation Process

The JIT compilation process typically consists of three stages: interpretation, profiling, and compilation.

1. Interpretation: Initially, the Execution Engine interprets the bytecode instructions one by one, executing them sequentially. This allows the JVM to quickly start executing the program without the need for compilation.
2. Profiling: During interpretation, the JVM collects profiling information about the program execution. This includes identifying frequently executed methods, hot spots, and code branches.
3. Compilation: Based on the profiling information, the JVM identifies hot methods or code segments that would benefit from compilation. These methods are then compiled into highly optimized native machine code. The compilation process may employ optimizations such as constant folding, dead code elimination, and loop unrolling.

Tiered Compilation

Modern JVMs often use a technique called tiered compilation to balance between interpreting and compiling code. In tiered compilation, the JVM uses multiple levels or tiers of compilation to optimize the execution of the program gradually.

- The first tier executes the program in interpreted mode to enable quick startup.
- The second tier performs quick compilations on hot methods to improve performance.
- The third tier performs more advanced and time-consuming optimizations on highly critical methods.

Code Caching

To further improve performance, the JVM also employs code caching. This means that the compiled native code is stored in memory for future use. By caching the compiled code, the JVM avoids the overhead of repeated compilation for frequently executed code segments.

Quiz

Check your knowledge answering some questions

04 | Quiz

Question 1/6

What is the role of JVM in the execution of Java programs?

- Converts Java bytecode to machine code
 - Allocates memory for Java objects
 - Provides a runtime environment to execute Java programs
-

Question 2/6

What is the first step in the JVM startup process?

- Loading the Java classes
 - Creating the Java objects
 - Initializing the JVM
-

Question 3/6

Which part of the JVM is responsible for loading Java classes?

- Garbage Collector
 - Class Loader
 - JIT Compiler
-

Question 4/6

What is the purpose of garbage collection in JVM?

- To optimize the execution of Java programs
 - To reclaim memory occupied by unused Java objects
 - To convert Java bytecode to machine code
-

Question 5/6

Which component of the JVM converts Java bytecode to machine code?

- Class Loader
 - Execution Engine
 - Garbage Collector
-

Question 6/6

What is the purpose of Just-In-Time (JIT) compilation in JVM?

- To allocate memory for Java objects
 - To improve the performance of Java programs
 - To load Java classes into memory
-

Submit

Conclusion

Congratulations!

Congratulations on completing this course! You have taken an important step in unlocking your full potential. Completing this course is not just about acquiring knowledge; it's about putting that knowledge into practice and making a positive impact on the world around you.



Share this course

Created with **LearningStudioAI**

v0.5.72