

This specification is not final and is subject to change. Use is subject to license terms.

API OTHER SPECIFICATIONS TOOL GUIDES

Java SE 22 & JDK 22  
DRAFT 22-internal-adhoc.gbierman.20231030

# Implicitly Declared Classes and Instance main Methods (Second Preview)

**Changes to the Java® Language Specification • Version 22-internal-adhoc.gbierman.20231030**

## Chapter 6: Names

### 6.3 Scope of a Declaration

## Chapter 7: Packages and Modules

### 7.3 Compilation Units

## Chapter 8: Classes

### 8.1 Class Declarations

## Chapter 12: Execution

### 12.1 Java Virtual Machine Startup

#### 12.1.1 Load the Initial Class or Interface ~~Test~~

#### 12.1.2 Link the Initial Class or Interface ~~Test~~: Verify, Prepare, (Optionally) Resolve

#### 12.1.3 Initialize the Initial Class or Interface ~~Test~~: Execute Initializers

#### 12.1.4 Invoke ~~Test.main~~ a main method

## Chapter 13: Binary Compatibility

### 13.1 The Form of a Binary

This document describes changes to the [Java Language Specification](#) to support *Implicitly Declared Classes and Instance main Methods*, which is a preview feature of Java SE 22. See [JEP 463](#) for an overview of the feature.

A companion document describes the changes needed to the [Java Virtual Machine Specification](#) to support Simple Compilation Units and Instance main Methods.

Changes are described with respect to existing sections of the JLS. New text is indicated **like this** and deleted text is indicated **like this**. Explanation and discussion, as needed, is set aside in grey boxes.

### Changelog:

2023-09-27: First draft of second preview. Main changes to first preview:

- Remove concept of an unnamed class in favour of a simple compilation unit implicitly declaring a class whose name is determined by the host system. This class name is never in scope and is not marked as synthetic.
- Removed the behavioral change that an inherited `static main` method would not be considered as a candidate method.
- Simplified the launcher semantics. Now a candidate method with a single formal parameter is always preferred over one with no parameters.

## Chapter 6: Names

### 6.3 Scope of a Declaration

The *scope* of a declaration is the region of the program within which the entity declared by the declaration can be referred to using a simple name, provided it is not shadowed (6.4.1 ↗).

A declaration is said to be *in scope* at a particular point in a program if and only if the declaration's scope includes that point.

The scope of the declaration of an observable top level package (7.4.3 ↗) is all observable compilation units associated with modules to which the package is uniquely visible (7.4.3 ↗).

The declaration of a package that is not observable is never in scope.

The declaration of a subpackage is never in scope.

The package `java` is always in scope.

The scope of a class or interface imported by a single-type-import declaration (7.5.1 ↗) or a type-import-on-demand declaration (7.5.2 ↗) is the module declaration (7.7 ↗) and all the class and interface declarations (8.1, 9.1 ↗) of the compilation unit in which the `import` declaration appears, as well as any annotations on the module declaration or package declaration of the compilation unit.

The scope of a member imported by a single-static-import declaration (7.5.3 ↗) or a static-import-on-demand declaration (7.5.4 ↗) is the module declaration and all the class and interface declarations of the compilation unit in which the `import` declaration appears, as well as any annotations on the module declaration or package declaration of the compilation unit.

The scope of a top level class or interface (7.6 ↗) declared in an ordinary compilation unit is all class and interface declarations in the package in which the top level class or interface is declared.

The implicit declaration of a top level class in a simple compilation unit (7.3) is never in scope.

The scope of a declaration of a member *m* declared in or inherited by a class or interface *C* (8.2 ↗, 9.2 ↗) is the entire body of *C*, including any nested class or interface declarations. If *C* is a record class, then the scope of *m* additionally includes the header of the record declaration of *C*.

*The rest of this section is unchanged.*

## Chapter 7: Packages and Modules

### 7.3 Compilation Units

*CompilationUnit* is the goal symbol (2.1 ↗) for the syntactic grammar (2.3 ↗) of Java programs. It is defined by the following production:

*CompilationUnit*:

*OrdinaryCompilationUnit*

*SimpleCompilationUnit*

*ModularCompilationUnit*

*OrdinaryCompilationUnit:*

`[PackageDeclaration] {ImportDeclaration} {TopLevelClassOrInterfaceDeclaration}`

*SimpleCompilationUnit:*

`{ImportDeclaration} {ClassMemberDeclarationNoMethod} MethodDeclaration  
{ClassMemberDeclaration}`

*ClassMemberDeclarationNoMethod:*

`FieldDeclaration`

`ClassDeclaration`

`InterfaceDeclaration`

`;-`

*ModularCompilationUnit:*

`{ImportDeclaration} ModuleDeclaration`

An *ordinary compilation unit* consists of three parts, each of which is optional:

- A `package` declaration (7.4 [↗](#)), giving the fully qualified name (6.7) of the package to which the compilation unit belongs.

A compilation unit that has no `package` declaration is part of an unnamed package (7.4.2 [↗](#)).

- `import` declarations (7.5 [↗](#)) that allow classes and interface from other packages, and `static` members of classes and interfaces, to be referred to using their simple names.
- Top level declarations of classes and interfaces (7.6 [↗](#)).

An *simple compilation unit* consists of:

- zero or more `import` declarations that allow classes and interface from other packages, and `static` members of classes and interfaces, to be referred to using their simple names.
- The declarations of the members (8.2 [↗](#)) of the implicitly declared top level class, at least one of which is a method declaration (8.4 [↗](#)).

*This means that the following compilation unit is unambiguously an ordinary compilation unit:*

```
import p.*;
class Test { ... }
```

*whereas the following is unambiguously a simple compilation unit:*

```
import p.*;
static void main(){ ... }
class Test { ... }
```

A simple compilation unit implicitly declares a class that satisfies the following properties:

- It is a top level class (7.6 [↗](#)).
- Its name is a valid identifier (3.8 [↗](#)) determined by the host system.

*In simple implementations of the Java SE Platform, where compilation units are stored in files, the name of this implicitly declared class would typically be the name of the file containing the simple compilation unit (7.3) minus any extension (such as `.java` or `.jav`).*

- It is not abstract (8.1.1.1 ↗).
- It is final (8.1.1.2 ↗).
- It is a member of an unnamed package (7.4.2 ↗) and has package access.
- Its direct superclass is Object (8.1.4 ↗).
- It does not have any direct superinterfaces (8.1.5 ↗).
- The body of the class contains every ClassMemberDeclaration (these are declarations of fields (8.3 ↗), methods (8.4 ↗), member classes (8.5 ↗), and member interfaces (9.1.1.3 ↗)) from the simple compilation unit. It is not possible for a simple compilation unit to declare an instance initializer (8.6 ↗), static initializer (8.7 ↗), or constructor (8.8 ↗).
- It has an implicitly declared default constructor (8.8.9 ↗).

All members of this class, including any implicitly declared members, are subject to the usual rules for member declarations in a class.

It is a compile-time error if this class does not declare a candidate main method (12.1.4).

*Note that an unnamed package may have multiple implicitly declared classes as members.*

A modular compilation unit consists of a `module` declaration (7.7 ↗), optionally preceded by `import` declarations. The `import` declarations allow classes and interfaces from packages in this module and other modules, as well as `static` members of classes and interfaces, to be referred to using their simple names within the `module` declaration.

Every compilation unit implicitly imports every `public` class or interface declared in the predefined package `java.lang`, as if the declaration `import java.lang.*;` appeared at the beginning of each compilation unit immediately after any `package` declaration. As a result, the names of all those classes and interfaces are available as simple names in every compilation unit.

The host system determines which compilation units are *observable*, except for the compilation units in the predefined package `java` and its subpackages `lang` and `io`, which are all always observable.

Each observable compilation unit may be *associated* with a module, as follows:

- The host system may determine that an observable ordinary compilation unit is associated with a module chosen by the host system, except for (i) the ordinary compilation units in the predefined package `java` and its subpackages `lang` and `io`, which are all associated with the `java.base` module, and (ii) any ordinary compilation unit in an unnamed package, which is associated with a module as specified in 7.4.2 ↗.
- The host system must determine that an observable modular compilation unit is associated with the module declared by the modular compilation unit.

*The observability of a compilation unit influences the observability of its package (7.4.3 ↗), while the association of an observable compilation unit with a module influences the observability of that module (7.7.6 ↗).*

When compiling the modular and ordinary compilation units associated with a module *M*, the host system must respect the dependencies specified in *M*'s declaration. Specifically, the host system must limit the ordinary compilation units that would otherwise be observable, to only

those that are *visible to M*. The ordinary compilation units that are visible to *M* are the observable ordinary compilation units associated with the modules that are *read by M*. The modules read by *M* are given by the result of *resolution*, as described in the `java.lang.module` package specification, with *M* as the only root module. The host system must perform resolution to determine the modules read by *M*; it is a compile-time error if resolution fails for any of the reasons described in the `java.lang.module` package specification.

*The readability relation is reflexive, so M reads itself, and thus all of the modular and ordinary compilation units associated with M are visible to M.*

*The modules read by M drive the packages that are uniquely visible to M (7.4.3 ↗), which in turn drives both the top level packages in scope and the meaning of package names for code in the modular and ordinary compilation units associated with M (6.3 ↗, 6.5.3 ↗, 6.5.5 ↗).*

*The rules above ensure that package and type names used in annotations in a modular compilation unit (in particular, annotations applied to the module declaration) are interpreted as if they appeared in an ordinary compilation unit associated with the module.*

Classes and interfaces declared in different ordinary compilation units can refer to each other, circularly. A Java compiler must arrange to compile all such classes and interfaces at the same time.

## Chapter 8: Classes

### 8.1 Class Declarations

A *class declaration* specifies a class.

There are three kinds of class declarations: *normal class declarations*, *enum declarations* (8.9 ↗), and *record declarations* (8.10 ↗).

*ClassDeclaration:*

*NormalClassDeclaration*

*EnumDeclaration*

*RecordDeclaration*

*NormalClassDeclaration:*

```
{ClassModifier} class TypeIdentifier [TypeParameters]
[ClassExtends] [ClassImplements] [ClassPermits] ClassBody
```

A class is also implicitly declared by **a simple compilation unit (7.3)**, a class instance creation expression (15.9.5 ↗) and an enum constant that ends with a class body (8.9.1 ↗).

The *TypeIdentifier* in a class declaration specifies the name of the class.

It is a compile-time error if a class has the same simple name as any of its enclosing classes or interfaces.

The scope and shadowing of a class declaration is specified in 6.3 ↗ and 6.4.1 ↗.

## Chapter 12: Execution

This chapter specifies activities that occur during execution of a program. It is organized around the life cycle of the Java Virtual Machine and of the classes, interfaces, and objects that form a program.

The Java Virtual Machine starts up by loading a specified class or interface, then invoking the a method `main` in this specified class or interface. Section 12.1 outlines the loading, linking, and initialization steps involved in executing `main`, as an introduction to the concepts in this chapter. Further sections specify the details of loading (12.2 ↗), linking (12.3 ↗), and initialization (12.4 ↗).

The chapter continues with a specification of the procedures for creation of new class instances (12.5 ↗); and finalization of class instances (12.6 ↗). It concludes by describing the unloading of classes (12.7 ↗) and the procedure followed when a program exits (12.8 ↗).

## 12.1 Java Virtual Machine Startup

The Java Virtual Machine starts execution by invoking the a method `main` of some specified class or interface. If this `main` method has a formal parameter, it is passed ~~passing it~~ a single argument which is an array of strings. ~~In the examples in this specification, this first class is typically called `Test`.~~

The precise semantics of Java Virtual Machine startup are given in Chapter 5 of *The Java Virtual Machine Specification, Java SE 22 Edition*. Here we present an overview of the process from the viewpoint of the Java programming language.

The manner in which the initial class or interface is specified to the Java Virtual Machine is beyond the scope of this specification, but it is typical, in host environments that use command lines, for the fully qualified name of the initial class or interface to be specified as a command line argument and for following command line arguments to be used as strings to be provided as the argument to the method `main`. If the original compilation unit was a simple compilation unit (7.3), then the name of the file that contained the compilation unit, minus any extension, is typically used to specify the name of the initial class or interface.

*For example, in a UNIX implementation, the command line:*

```
java Test reboot Bob Dot Enzo
```

*will typically start a Java Virtual Machine by invoking method `main` of class `Test` (a class in an unnamed package), passing it an argument array containing the four strings "reboot", "Bob", "Dot", and "Enzo".*

Whereas if the file `HelloWorld.java` contained the following simple compilation unit:

```
void main() {
    System.out.println("Hello, World!");
}
```

which has been compiled, then the command line:

```
java HelloWorld
```

will typically start a Java Virtual Machine by invoking the `main` method of the implicitly declared class (7.3) producing the output:

```
Hello, World!
```

We now outline the steps the Java Virtual Machine may take to execute ~~`Test`~~ the initial class or interface, as an example of the loading, linking, and initialization processes that are described further in later sections.

### 12.1.1 Load the Initial Class or Interface ~~`Test`~~

The initial attempt to execute the a method `main` of the initial class or interface ~~`Test`~~ discovers

that the class `Test` `it` is not loaded - that is, that the Java Virtual Machine does not currently contain a binary representation for this class or interface. The Java Virtual Machine then uses a class loader to attempt to find such a binary representation. If this process fails, then an error is thrown. This loading process is described further in [12.2](#).

### 12.1.2 Link the Initial Class or Interface `Test`: Verify, Prepare, (Optionally) Resolve

After the class or interface `Test` is loaded, it must be initialized before a method `main` can be invoked. And `Test`, like all classes and interfaces, it must be linked before it is initialized. Linking involves verification, preparation, and (optionally) resolution. Linking is described further in [12.3](#).

Verification checks that the loaded representation of the class or interface `Test` is well-formed, with a proper symbol table. Verification also checks that the code that implements the class or interface `Test` obeys the semantic requirements of the Java programming language and the Java Virtual Machine. If a problem is detected during verification, then an error is thrown. Verification is described further in [12.3.1](#).

Preparation involves allocation of static storage and any data structures that are used internally by the implementation of the Java Virtual Machine, such as method tables. Preparation is described further in [12.3.2](#).

Resolution is the process of checking symbolic references from the class or interface `Test` to other classes and interfaces, by loading the other classes and interfaces that are mentioned and checking that the references are correct.

The resolution step is optional at the time of initial linkage. An implementation may resolve symbolic references from a class or interface that is being linked very early, even to the point of resolving all symbolic references from the classes and interfaces that are further referenced, recursively. (This resolution may result in errors from these further loading and linking steps.) This implementation choice represents one extreme and is similar to the kind of "static" linkage that has been done for many years in simple implementations of the C language. (In these implementations, a compiled program is typically represented as an "a.out" file that contains a fully-linked version of the program, including completely resolved links to library routines used by the program. Copies of these library routines are included in the "a.out" file.)

An implementation may instead choose to resolve a symbolic reference only when it is actively used; consistent use of this strategy for all symbolic references would represent the "laziest" form of resolution. In this case, if the class or interface `Test` had several symbolic references to another class, then the references might be resolved one at a time, as they are used, or perhaps not at all, if these references were never used during execution of the program.

The only requirement on when resolution is performed is that any errors detected during resolution must be thrown at a point in the program where some action is taken by the program that might, directly or indirectly, require linkage to the class or interface involved in the error. Using the "static" example implementation choice described above, loading and linkage errors could occur before the program is executed if they involved a class or interface mentioned in the initial class or interface `Test` or any of the further, recursively referenced, classes and interfaces. In a system that implemented the "laziest" resolution, these errors would be thrown only when an incorrect symbolic reference is actively used.

The resolution process is described further in [12.3.3](#).

### 12.1.3 Initialize the Initial Class or Interface `Test`: Execute Initializers

In our continuing example, the Java Virtual Machine is still trying to execute **the a** method `main` of **the initial** class **or interface** `Test`. This is permitted only if the class has been initialized (12.4.1 ↗).

Initialization consists of execution of any class variable initializers and static initializers of the **initial** class **or interface** `Test`, in textual order. But before **it** `Test` can be initialized, its direct superclass must be initialized, as well as the direct superclass of its direct superclass, and so on, recursively. In the simplest case, **the initial class or interface** `Test` has `Object` as its implicit direct superclass; if class `Object` has not yet been initialized, then it must be initialized before **the initial class or interface** `Test` is initialized. Class `Object` has no superclass, so the recursion terminates here.

If **the initial** class **or interface** `Test` has another class `Super` as its superclass, then `Super` must be initialized before **the initial class or interface** `Test`. This requires loading, verifying, and preparing `Super` if this has not already been done and, depending on the implementation, may also involve resolving the symbolic references from `Super` and so on, recursively.

Initialization may thus cause loading, linking, and initialization errors, including such errors involving other classes and interfaces.

The initialization process is described further in 12.4 ↗.

#### 12.1.4 Invoke `Test.main` a main method

Finally, after completion of the initialization for **the initial** class **or interface** `Test` (during which other consequential loading, linking, and initializing may have occurred), **the a main** method `main` of **the initial class or interface** `Test` is invoked.

~~The method `main` must be declared `public`, `static`, and `void`. It must specify a formal parameter (8.4.1 ↗) whose declared type is array of `String`. Therefore, either of the following declarations is acceptable:~~

```
public static void main(String[] args)
public static void main(String... args)
```

**A method of the initial class or interface is a *candidate* if it is named `main` and one of the following applies:**

- **It has a single formal parameter (8.4.1 ↗) whose declared type is an array of `String` and a `void` result, and is with `public`, `protected` or package access.**
- **It has no formal parameters and a `void` result, and is with `public`, `protected` or package access.**

*Note that a candidate `main` method may be either a `static` or an instance method. A candidate `main` method may also have a `throws` clause (8.4.6 ↗).*

*The permitted signature of a `main` method expanded significantly in Java SE 22. Prior to that, it had to be `public static` and have a single formal parameter; the only variation possible was `String[]` versus `String...` for the type of the single formal parameter. In Java 22 and above, `main` can have one of twelve possible signatures: six with a single formal parameter and six with no formal parameters. This number increases to 18 if `String[]` is distinguished from `String...` in the type of the single formal parameter.*

*Note that it is not a compile-time error if the initial class or interface counts more than one candidate `main` method among its members.*



*The presence of a `main` method in a class or interface may not be immediately apparent because a non-static `main` method may be inherited. For example, a default method in an interface is an instance method (9.4.2), so may be a candidate when inherited by a class that implements the interface. Development tools are encouraged to highlight when a class or interface has a member `main` method that could serve as the start of the program.*

A `main` method of the initial class or interface is invoked, as if by application of the following rules:

- If there is a candidate method with a single formal parameter then this method is invoked. If the candidate method is `static` then it is invoked directly, passing the argument array (12.1). If the candidate method is an instance method then it is invoked, passing the argument array, on an instance of the initial class created by using a constructor with no formal parameters and either `public`, `protected`, or package access.
- Otherwise, if there is a candidate method with no formal parameters then this method is invoked. If the candidate method is `static` then it is invoked directly. If the candidate method is an instance method then it is invoked on an instance of the initial class created by using a constructor with no formal parameters and either `public`, `protected`, or package access.

*Note that it is not possible for a class to have both a `static` and an instance method with the same signature (8.4.2.2).*

*The behavior of an implementation if there is no candidate method to invoke, or if there is no suitable constructor in the initial class when invoking an instance candidate method, is beyond the scope of this specification.*

## Chapter 13: Binary Compatibility

### 13.1 The Form of a Binary

Programs must be compiled either into the `class` file format specified by *The Java Virtual Machine Specification, Java SE 22 Edition*, or into a representation that can be mapped into that format by a class loader written in the Java programming language.

A `class` file corresponding to a class or interface declaration must have certain properties. A number of these properties are specifically chosen to support source code transformations that preserve binary compatibility. The required properties are:

1. The class or interface must be named by its *binary name*, which must meet the following constraints:
  - The binary name of a top level class or interface (7.6.2) is its canonical name (6.7).
  - The binary name of a member class or interface (8.5.2, 9.5.2) consists of the binary name of its immediately enclosing class or interface, followed by `$`, followed by the simple name of the member.
  - The binary name of a local class or interface (14.3.2) consists of the binary name of its immediately enclosing class or interface, followed by `$`, followed by a non-empty sequence of digits, followed by the simple name of the local class.
  - The binary name of an anonymous class (15.9.5.2) consists of the binary name of its immediately enclosing class or interface, followed by `$`, followed by a non-empty sequence of digits.

- The binary name of a type variable declared by a generic class or interface (8.1.2 ↗, 9.1.2 ↗) is the binary name of its immediately enclosing class or interface, followed by \$, followed by the simple name of the type variable.
  - The binary name of a type variable declared by a generic method (8.4.4 ↗) is the binary name of the class or interface declaring the method, followed by \$, followed by the descriptor of the method (JVMS §4.3.3), followed by \$, followed by the simple name of the type variable.
  - The binary name of a type variable declared by a generic constructor (8.8.4 ↗) is the binary name of the class declaring the constructor, followed by \$, followed by the descriptor of the constructor (JVMS §4.3.3), followed by \$, followed by the simple name of the type variable.
2. A reference to another class or interface must be symbolic, using the binary name of the class or interface.
  3. A reference to a field that is a constant variable (4.12.4 ↗) must be resolved at compile time to the value  $V$  denoted by the constant variable's initializer.

If such a field is `static`, then no reference to the field should be present in the code in a binary file, including the class or interface which declared the field. Such a field must always appear to have been initialized (12.4.2 ↗); the default initial value for the field (if different than  $V$ ) must never be observed.

If such a field is `non-static`, then no reference to the field should be present in the code in a binary file, except in the class containing the field. (It will be a class rather than an interface, since an interface has only `static` fields.) The class should have code to set the field's value to  $V$  during instance creation (12.5 ↗).

4. Given a legal expression denoting a field access in a class  $C$ , referencing a field named  $f$  that is not a constant variable and is declared in a (possibly distinct) class or interface  $D$ , we define the *qualifying class or interface of the field reference* as follows:
  - If the expression is referenced by a simple name, then if  $f$  is a member of the current class or interface,  $C$ , then let  $Q$  be  $C$ . Otherwise, let  $Q$  be the innermost lexically enclosing class or interface declaration of which  $f$  is a member. In either case,  $Q$  is the qualifying class or interface of the reference.
  - If the reference is of the form  $TypeName.f$ , where  $TypeName$  denotes a class or interface, then the class or interface denoted by  $TypeName$  is the qualifying class or interface of the reference.
  - If the expression is of the form  $ExpressionName.f$  or  $Primary.f$ , then:
    - If the compile-time type of  $ExpressionName$  or  $Primary$  is an intersection type  $V_1 \& \dots \& V_n$  (4.9 ↗), then the qualifying class or interface of the reference is the erasure (4.6 ↗) of  $V_1$ .
    - Otherwise, the erasure of the compile-time type of  $ExpressionName$  or  $Primary$  is the qualifying class or interface of the reference.
  - If the expression is of the form  $super.f$ , then the superclass of  $C$  is the qualifying class or interface of the reference.
  - If the expression is of the form  $TypeName.super.f$ , then the superclass of the class denoted by  $TypeName$  is the qualifying class or interface of the reference.

The reference to  $f$  must be compiled into a symbolic reference to the qualifying class or interface of the reference, plus the simple name of the field,  $f$ .

The reference must also include a symbolic reference to the erasure of the declared type of the field, so that the verifier can check that the type is as expected.

5. Given a method invocation expression or a method reference expression in a class or interface  $C$ , referencing a method named  $m$  declared (or implicitly declared (9.2 ↗)) in a (possibly distinct) class or interface  $D$ , we define the *qualifying class or interface of the method invocation* as follows:
  - If  $D$  is `Object` then the qualifying class or interface of the method invocation is `Object`.
  - Otherwise:
    - If the method is referenced by a simple name, then if  $m$  is a member of the current class or interface  $C$ , let  $Q$  be  $C$ ; otherwise, let  $Q$  be the innermost lexically enclosing class or interface declaration of which  $m$  is a member. In either case,  $Q$  is the qualifying class or interface of the method invocation.
    - If the expression is of the form  $TypeName.m$  or  $ReferenceType::m$ , then the class or interface denoted by  $TypeName$ , or the erasure of  $ReferenceType$ , is the qualifying class or interface of the method invocation.
    - If the expression is of the form  $ExpressionName.m$  or  $Primary.m$  or  $ExpressionName::m$  or  $Primary::m$ , then:
      - If the compile-time type of  $ExpressionName$  or  $Primary$  is an intersection type  $V_1 \& \dots \& V_n$ , then the qualifying class or interface of the method invocation is the erasure of  $V_1$ .
      - Otherwise, the erasure of the compile-time type of  $ExpressionName$  or  $Primary$  is the qualifying class or interface of the method invocation.
    - If the expression is of the form  $super.m$  or  $super::m$ , then the superclass of  $C$  is the qualifying class or interface of the method invocation.
    - If the expression is of the form  $TypeName.super.m$  or  $TypeName.super::m$ , then if  $TypeName$  denotes a class  $X$ , the superclass of  $X$  is the qualifying class or interface of the method invocation; if  $TypeName$  denotes an interface  $X$ ,  $X$  is the qualifying class or interface of the method invocation.

A reference to a method must be resolved at compile time to a symbolic reference to the qualifying class or interface of the method invocation, plus the erasure of the declared signature (8.4.2 ↗) of the method. The signature of a method must include all of the following as determined by 15.12.3 ↗:

- The simple name of the method
- The number of parameters to the method
- A symbolic reference to the type of each parameter

A reference to a method must also include either a symbolic reference to the erasure of the return type of the denoted method or an indication that the denoted method is declared `void` and does not return a value.

6. Given a class instance creation expression (15.9 ↗) or an explicit constructor invocation

statement (8.8.7.1 ↗) or a method reference expression of the form *ClassType* :: *new* (15.13 ↗) in a class or interface *C*, referencing a constructor *m* declared in a (possibly distinct) class or interface *D*, we define the *qualifying class of the constructor invocation* as follows:

- If the expression is of the form *new D(...)* or *ExpressionName.new D(...)* or *Primary.new D(...)* or *D :: new*, then the qualifying class of the constructor invocation is *D*.
- If the expression is of the form *new D(...){...}* or *ExpressionName.new D(...){...}* or *Primary.new D(...){...}*, then the qualifying class of the constructor invocation is the anonymous class declared by the expression.
- If the expression is of the form *super(...)* or *ExpressionName.super(...)* or *Primary.super(...)*, then the qualifying class of the constructor invocation is the direct superclass of *C*.
- If the expression is of the form *this(...)*, then the qualifying class of the constructor invocation is *C*.

A reference to a constructor must be resolved at compile time to a symbolic reference to the qualifying class of the constructor invocation, plus the declared signature of the constructor (8.8.2 ↗). The signature of a constructor must include both:

- The number of parameters of the constructor
- A symbolic reference to the type of each formal parameter

A binary representation for a class or interface must also contain all of the following:

1. If it is a class and is not `Object`, then a symbolic reference to the direct superclass of this class.
2. A symbolic reference to each direct superinterface, if any.
3. A specification of each field declared in the class or interface, given as the simple name of the field and a symbolic reference to the erasure of the type of the field.
4. If it is a class, then the erased signature of each constructor, as described above.
5. For each method declared in the class or interface (excluding, for an interface, its implicitly declared methods (9.2 ↗)), its erased signature and return type, as described above.
6. The code needed to implement the class or interface:
  - For an interface, code for the field initializers and the implementation of each method with a block body (9.4.3 ↗).
  - For a class, code for the field initializers, the instance and static initializers, the implementation of each method with a block body (8.4.7 ↗), and the implementation of each constructor.
7. Every class or interface must contain sufficient information to recover its canonical name (6.7).
8. Every member class or interface must have sufficient information to recover its source-level access modifier (6.6 ↗).
9. Every nested class or interface must have a symbolic reference to its immediately

enclosing class or interface (8.1.3).

10. Every class or interface must contain symbolic references to all of its member classes and interfaces (8.5 ↗, 9.5 ↗), and to all other nested classes and interfaces declared within its body.
11. A construct emitted by a Java compiler must be marked as *synthetic* if it does not correspond to a construct declared explicitly or implicitly in source code, unless the emitted construct is a class initialization method (JVMS §2.9).
12. A construct emitted by a Java compiler must be marked as *mandated* if it corresponds to a formal parameter declared implicitly in source code (8.8.1 ↗, 8.8.9 ↗, 8.9.3 ↗, 15.9.5.1 ↗).

The following formal parameters are declared implicitly in source code:

- The first formal parameter of a constructor of a non-private inner member class (8.8.1 ↗, 8.8.9 ↗).
- The first formal parameter of an anonymous constructor of an anonymous class whose superclass is an inner class (not in a static context) (15.9.5.1 ↗).
- The formal parameter *name* of the *valueOf* method which is implicitly declared in an enum class (8.9.3 ↗).
- The formal parameters of a compact constructor of a record class (8.10.4 ↗).

For reference, the following constructs are declared implicitly in source code, but are not marked as mandated because only formal parameters and modules can be so marked in a `class` file (JVMS §4.7.24, JVMS §4.7.25):

- Default constructors of normal and enum classes (8.8.9 ↗, 8.9.2 ↗)
- Canonical constructors of record classes (8.10.4 ↗)
- Anonymous constructors (15.9.5.1 ↗)
- The *values* and *valueOf* methods of enum classes (8.9.3 ↗)
- Certain *public* fields of enum classes (8.9.3 ↗)
- Certain *private* fields and *public* methods of record classes (8.10.3 ↗)
- Certain *public* methods of interfaces (9.2 ↗)
- Container annotations (9.7.5 ↗)
- Classes implicitly declared in a simple compilation unit (7.3).

A `class` file corresponding to a module declaration must have the properties of a `class` file for a class whose binary name is `module-info` and which has no superclass, no superinterfaces, no fields, and no methods. In addition, the binary representation of the module must contain all of the following:

- A specification of the name of the module, given as a symbolic reference to the name indicated after `module`. Also, the specification must include whether the module is normal or open (7.7 ↗).
- A specification of each dependence denoted by a `requires` directive, given as a symbolic reference to the name of the module indicated by the directive (7.7.1 ↗). Also, the specification must include whether the dependence is `transitive` and whether the dependence is `static`.
- A specification of each package denoted by an `exports` or `opens` directive, given as a

symbolic reference to the name of the package indicated by the directive (7.7.2 ↗). Also, if the directive was qualified, the specification must give symbolic references to the names of the modules indicated by the directive's `to` clause.

- A specification of each service denoted by a `uses` directive, given as a symbolic reference to the name of the class or interface indicated by the directive (7.7.3 ↗).
- A specification of the service providers denoted by a `provides` directive, given as symbolic references to the names of the classes and interfaces indicated by the directive's `with` clause (7.7.4 ↗). Also, the specification must give a symbolic reference to the name of the class or interface indicated as the service by the directive.

The following sections discuss changes that may be made to class and interface declarations without breaking compatibility with pre-existing binaries. Under the translation requirements given above, the Java Virtual Machine and its `class` file format support these changes. Any other valid binary format, such as a compressed or encrypted representation that is mapped back into `class` files by a class loader under the above requirements, will necessarily support these changes as well.

---

*Copyright © 1993, 2023, Oracle and/or its affiliates, 500 Oracle Parkway, Redwood Shores, CA 94065 USA. All rights reserved. Use is subject to [license terms](#) and the [documentation redistribution policy](#).*

**DRAFT 22-internal-adhoc.gbierman.20231030**