Java Platform, Standard Edition Java Language Updates





Java Platform, Standard Edition Java Language Updates, Release 16

F35146-01

Copyright © 2017, 2021, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface
Audience
Documentation Accessibility
Related Documents
Conventions

1 Java Language Changes

Java Language Updates for Java SE 16	1-1
Java Language Updates for Java SE 15	1-2
Java Language Changes for Java SE 14	1-3
Java Language Changes for Java SE 13	1-3
Java Language Changes for Java SE 12	1-4
Java Language Changes for Java SE 11	1-4
Java Language Changes for Java SE 10	1-4
Java Language Changes for Java SE 9	1-4
More Concise try-with-resources Statements	1-5
@SafeVarargs Annotation Allowed on Private Instance Methods	1-6
Diamond Syntax and Anonymous Inner Classes	1-6
Underscore Character Not Legal Name	1-6
Support for Private Interface Methods	1-6

2 Preview Features

3 Sealed Classes

4 Pattern Matching for instanceof



5	Record Classes
6	Switch Expressions
7	Text Blocks
8	Local Variable Type Inference



Preface

This guide describes the updated language features in Java SE 9 and subsequent releases.

Audience

This document is for Java developers.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

See JDK 16 Documentation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.



1

Java Language Changes

This section summarizes the updated language features in Java SE 9 and subsequent releases.

Java Language Updates for Java SE 16

Feature	Description	JEP
Sealed Classes	Preview feature from Java SE 15 re-previewed for this release. It has been enhanced with several refinements, including more strict checking of narrowing reference conversions with respect to sealed type hierarchies.	JEP 397: Sealed Classes (Second Preview)
	A sealed class or interface restricts which classes or interfaces can extend or implement it.	
Record Classes	First previewed in Java SE 14, this feature is permanent in this release. This means that it can be used in any program compiled for Java SE 16 without enabling preview features.	JEP 395: Records
	In this release, inner classes may declare members that are either explicitly or implicitly static. This includes record class members, which are implicitly static.	
	A record is a class that acts as transparent carrier for immutable data.	



Feature	Description	JEP
Pattern Matching for instanceof	First previewed in Java SE 14, this feature is permanent in this release. This means that it can be used in any program compiled for Java SE 16 without enabling preview features.	JEP 394: Pattern Matching for instanceof
	In this release, pattern variables are no longer implicitly final, and it's a compile-time error if a pattern instanceof expression compares an expression of type S with a pattern of type T, where S is a subtype of T.	
	Pattern matching allows common logic in a program, namely the conditional extraction of components from objects, to be expressed more concisely and safely.	

Java Language Updates for Java SE 15

Feature	Description	JEP
Sealed Classes	Introduced as a preview feature for this release. A sealed class or interface restricts which classes or interfaces can extend or implement it.	JEP 360: Sealed Classes (Preview)
Record Classes	Preview feature from Java SE 14 re-previewed for this release. It has been enhanced with support for local records.	JEP 384: Records (Second Preview)
	A record is a class that acts as transparent carrier for immutable data.	
Pattern Matching for instanceof	Preview feature from Java SE 14 re-previewed for this release. It is unchanged between Java SE 14 and this release.	JEP 375: Pattern Matching for instanceof (Second Preview)
	Pattern matching allows common logic in a program, namely the conditional extraction of components from objects, to be expressed more concisely and safely.	
Text Blocks	First previewed in Java SE 13, this feature is permanent	JEP 378: Text Blocks
See also Programmer's Guide to Text Blocks	in this release. This means that it can be used in any program compiled for Java SE 15 without enabling preview features.	
	A text block is a multiline string literal that avoids the need for most escape sequences, automatically formats the string in a predictable way, and gives the developer control over the format when desired.	



Java Language Changes for Java SE 14

Feature	Description	JEP
Pattern Matching for the instanceof Operator	Introduced as a preview feature for this release. Pattern matching allows common logic in a program, namely the conditional extraction of components from objects, to be expressed more concisely and safely.	JEP 305: Pattern Matching for instanceof (Preview)JEP 305: Pattern Matching for instanceof (Preview)
Records	Introduced as a preview feature for this release. Records provide a compact syntax for declaring classes which are transparent holders for shallowly immutable data.	JEP 359: Records (Preview)
Text Blocks See also Programmer's Guide to Text Blocks	Preview feature from Java SE 13 re-previewed for this release. It has been enhanced with support for more escape sequences. A text block is a multiline string literal that avoids the need for most escape sequences, automatically formats the string in a predictable way, and gives the developer control over the format when desired.	JEP 375: Pattern Matching for instanceof (Second Preview)
Switch Expressions	First previewed in Java SE 12, this feature is permanent in this release. This means that it can be used in any program compiled for Java SE 14 without needing to enable preview features. This feature extends switch so it can be used as either a statement or an expression, and so that both forms can use either traditional case: labels (with fall through) or new case> labels (with no fall through), with a further new statement for yielding a value from a switch expression.	JEP 361: Switch Expressions (Standard)

Java Language Changes for Java SE 13

Feature	Description	JEP
Text Blocks, see Programmer's Guide to Text Blocks	Introduced as a preview feature for this release. A text block is a multi-line string literal that avoids the need for most escape sequences, automatically formats the string in a predictable way, and gives the developer control over format when desired.	JEP 355: Text Blocks (Preview)
Switch Expressions	Preview feature from Java SE 12 re-previewed for this release. It has been enhanced with one change: To specify the value of a switch expression, use the new yield statement instead of the break statement.	JEP 354: Switch Expressions (Second Preview)
	This feature extends switch so it can be used as either a statement or an expression, and so that both forms can use either traditional case: labels (with fall through) or new case> labels (with no fall through), with a further new statement for yielding a value from a switch expression.	



Java Language Changes for Java SE 12

Feature	Description	JEP
Switch Expressions	Introduced as a preview feature for this release.	JEP 325: Switch
	This feature extends the switch statement so that it can be used as either a statement or an expression, and that both forms can use either a "traditional" or "simplified" scoping and control flow behavior.	Expressions (Preview)

Java Language Changes for Java SE 11

Feature	Description	JEP
Local Variable Type Inference See also Local Variable Type Inference: Style Guidelines	Introduced in Java SE 10. In this release, it has been enhanced with support for allowing var to be used when declaring the formal parameters of implicitly typed lambda expressions. Local-Variable Type Inference extends type inference to declarations of local variables with initializers.	 JEP 286: Local- Variable Type Inference JEP 323: Local- Variable Syntax for Lambda Parameters

Java Language Changes for Java SE 10

Feature	Description	JEP
Local Variable Type Inference	Introduced in this release. Local-Variable Type Inference extends type inference to	JEP 286: Local-Variable Type Inference
See also Local Variable Type Inference: Style Guidelines	declarations of local variables with initializers.	

Java Language Changes for Java SE 9

Feature	Description	JEP
Java Platform module system, see Project Jigsaw on OpenJDK.	Introduced in this release. The Java Platform module system introduces a new kind of Java programing component, the module, which is a named, self-describing collection of code and data. Its code is organized as a set of packages containing types, that is, Java classes and interfaces; its data includes resources and other kinds of static information. Modules can either export or encapsulate packages, and they express dependencies on other modules explicitly.	Java Platform Module System (JSR 376) JEP 261: Module System JEP 200: The Modular JDK JEP 220: Modular Run-Time Images JEP 260: Encapsulate Most Internal APIs



Feature		Description	JEP
Small language enhancements		Introduced in Java SE 7 as Project Coin. It has been enhanced with a few amendments.	JEP 213: Milling Project Coin
(Pr	oject Coin): More Concise try-with-resources		JSR 334: Small Enhancements to the Java Programming
	Statements		Language
•	@SafeVarargs Annotation Allowed on Private Instance Methods		
•	Diamond Syntax and Anonymous Inner Classes		
•	Underscore Character Not Legal Name		
•	Support for Private Interface Methods		

More Concise try-with-resources Statements

If you already have a resource as a final or effectively final variable, you can use that variable in a try-with-resources statement without declaring a new variable. An "effectively final" variable is one whose value is never changed after it is initialized.

For example, you declared these two resources:

```
// A final resource
final Resource resource1 = new Resource("resource1");
// An effectively final resource
Resource resource2 = new Resource("resource2");
```

In Java SE 7 or 8, you would declare new variables, like this:

```
try (Resource r1 = resource1;
    Resource r2 = resource2) {
    ...
}
```

In Java SE 9, you don't need to declare r1 and r2:

There is a more complete description of the try-with-resources statement in The Java Tutorials (Java SE 8 and earlier).

@SafeVarargs Annotation Allowed on Private Instance Methods

The @SafeVarargs annotation is allowed on private instance methods. It can be applied only to methods that cannot be overridden. These include static methods, final instance methods, and, new in Java SE 9, private instance methods.

Diamond Syntax and Anonymous Inner Classes

You can use diamond syntax in conjunction with anonymous inner classes. Types that can be written in a Java program, such as int or String, are called denotable types. The compiler-internal types that cannot be written in a Java program are called non-denotable types.

Non-denotable types can occur as the result of the inference used by the diamond operator. Because the inferred type using diamond with an anonymous class constructor could be outside of the set of types supported by the signature attribute in class files, using the diamond with anonymous classes was not allowed in Java SE 7.

Underscore Character Not Legal Name

If you use the underscore character ("_") as an identifier, your source code can no longer be compiled.

Support for Private Interface Methods

Private interface methods are supported. This support allows nonabstract methods of an interface to share code between them.



2

Preview Features

A preview feature is a new feature whose design, specification, and implementation are complete, but which is not permanent, which means that the feature may exist in a different form or not at all in future JDK releases.

Introducing a feature as a preview feature in a mainline JDK release enables the largest developer audience possible to try the feature out in the real world and provide feedback. In addition, tool vendors are encouraged to build support for the feature before Java developers use it in production. Developer feedback helps determine whether the feature has any design mistakes, which includes hard technical errors (such as a flaw in the type system), soft usability problems (such as a surprising interaction with an older feature), or poor architectural choices (such as one that forecloses on directions for future features). Through this feedback, the feature's strengths and weaknesses are evaluated to determine if the feature has a long-term role in the Java SE Platform, and if so, whether it needs refinement. Consequently, the feature may be granted final and permanent status (with or without refinements), or undergo a further preview period (with or without refinements), or else be removed.

Every preview feature is described by a JDK Enhancement Proposal (JEP) that defines its scope and sketches its design. For example, JEP 325 describes the JDK 12 preview feature for switch expressions. For background information about the role and lifecycle of preview features, see JEP 12.

Using Preview Features

To use preview language features in your programs, you must explicitly enable them in the compiler and the runtime system. If not, you'll receive an error message that states that your code is using a preview feature and preview features are disabled by default.

To compile source code with javac that uses preview features from JDK release n, use javac from JDK release n with the --enable-preview command-line option in conjunction with either the --release n or -source n command-line option.

For example, suppose you have an application named MyApp. java that uses the JDK 12 preview language feature switch expressions. Compile this with JDK 12 as follows:

javac --enable-preview --release 12 MyApp.java



Note:

When you compile an application that uses preview features, you'll receive a warning message similar to the following:

Note: MyApp.java uses preview language features. Note: Recompile with -Xlint:preview for details

Remember that preview features are subject to change and are intended to provoke feedback.

To run an application that uses preview features from JDK release n, use java from JDK release n with the --enable-preview option. To continue the previous example, to run MyApp, run java from JDK 12 as follows:

java --enable-preview MyApp

Note:

Code that uses preview features from an older release of the Java SE Platform will not necessarily compile or run on a newer release.

The tools jshell and javadoc also support the --enable-preview command-line option.

Sending Feedback

You can provide feedback on preview features, or anything else about the Java SE Platform, as follows:

- If you find any bugs, then submit them at Java Bug Database.
- If you want to provide substantive feedback on the usability of a preview feature, then post it on the OpenJDK mailing list where the feature is being discussed. To find the mailing list of a particular feature, see the feature's JEP page and look for the label *Discussion*. For example, on the page JEP 325: Switch Expressions (Preview), you'll find "*Discussion* amber dash dev at openjdk dot java dot net" near the top of the page.
- If you are working on an open source project, then see Quality Outreach on the OpenJDK Wiki.



Sealed Classes

Sealed classes and interfaces restrict which other classes or interfaces may extend or implement them.



This is a preview feature, which is a feature whose design, specification, and implementation are complete, but is not permanent, which means that the feature may exist in a different form or not at all in future Java SE releases. To compile and run code that contains preview features, you must specify additional command-line options. See Preview Features.

For background information about sealed classes and interfaces, see JEP 397.

One of the primary purposes of inheritance is code reuse: When you want to create a new class and there is already a class that includes some of the code that you want, you can derive your new class from the existing class. In doing this, you can reuse the fields and methods of the existing class without having to write (and debug) them yourself.

However, what if you want to model the various possibilities that exist in a domain by defining its entities and determining how these entities should relate to each other? For example, you're working on a graphics library. You want to determine how your library should handle common geometric primitives like circles and squares. You've created a Shape class that these geometric primitives can extend. However, you're not interested in allowing any arbitrary class to extend Shape; you don't want clients of your library declaring any further primitives. By sealing a class, you can specify which classes are permitted to extend it and prevent any other arbitrary class from doing so.

Defining Sealed Classes

To seal a class, add the sealed modifier to its declaration. Then, after any extends and implements clauses, add the permits clause. This clause specifies the classes that may extend the sealed class.

For example, the following declaration of Shape specifies three permitted subclasses, Circle, Square, and Rectangle:

Figure 3-1 Shape.java

```
public sealed class Shape
    permits Circle, Square, Rectangle {
}
```



Define the following three permitted subclasses, Circle, Square, and Rectangle, in the same module or in the same package as the sealed class:

Figure 3-2 Circle.java

```
public final class Circle extends Shape {
    public float radius;
}
```

Figure 3-3 Square.java

```
public non-sealed class Square extends Shape {
   public double side;
}
```

Figure 3-4 Rectangle.java

```
public sealed class Rectangle extends Shape permits FilledRectangle {
    public double length, width;
}
```

Rectangle has a further subclass, FilledRectangle:

Figure 3-5 FilledRectangle.java

```
public final class FilledRectangle extends Rectangle {
    public int red, green, blue;
}
```

Alternatively, you can define permitted subclasses in the same file as the sealed class. If you do so, then you can omit the permits clause:

```
package com.example.geometry;

public sealed class Figure
    // The permits clause has been omitted
    // as its permitted classes have been
    // defined in the same file.

{ }

final class Circle extends Figure {
    float radius;
}

non-sealed class Square extends Figure {
    float side;
```



```
}
sealed class Rectangle extends Figure {
    float length, width;
}
final class FilledRectangle extends Rectangle {
    int red, green, blue;
}
```

Constraints on Permitted Subclasses

Permitted subclasses have the following constraints:

They must be accessible by the sealed class at compile time.

For example, to compile Shape.java, the compiler must be able to access all of the permitted classes of Shape: Circle.java, Square.java, and Rectangle.java. In addition, because Rectangle is a sealed class, the compiler also needs access to FilledRectangle.java.

- They must directly extend the sealed class.
- They must have exactly one of the following modifiers to describe how it continues the sealing initiated by its superclass:
 - final: Cannot be extended further
 - sealed: Can only be extended by its permitted subclasses
 - non-sealed: Can be extended by unknown subclasses; a sealed class cannot prevent its permitted subclasses from doing this

For example, the permitted subclasses of Shape demonstrate each of these three modifiers: Circle is final while Rectangle is sealed and Square is non-sealed.

 They must be in the same module as the sealed class (if the sealed class is in a named module) or in the same package (if the sealed class is in the unnamed module, as in the Shape.java example).

For example, in the following declaration of <code>com.example.graphics.Shape</code>, its permitted subclasses are all in different packages. This example will compile only if <code>Shape</code> and all of its permitted subclasses are in the same named module.

Defining Sealed Interfaces

Like sealed classes, to seal an interface, add the sealed modifier to its declaration. Then, after any extends clause, add the permits clause, which specifies the classes that can implement the sealed interface and the interfaces that can extend the sealed interface.



The following example declares a sealed interface named Expr. Only the classes ConstantExpr, PlusExpr, TimesExpr, and NegExpr may implement it:

```
package com.example.expressions;
public class TestExpressions {
  public static void main(String[] args) {
    // (6 + 7) * -8
    System.out.println(
      new TimesExpr(
        new PlusExpr(new ConstantExpr(6), new ConstantExpr(7)),
        new NegExpr(new ConstantExpr(8))
      ).eval());
   }
}
sealed interface Expr
    permits ConstantExpr, PlusExpr, TimesExpr, NegExpr {
    public int eval();
}
final class ConstantExpr implements Expr {
    ConstantExpr(int i) { this.i = i; }
    public int eval() { return i; }
final class PlusExpr implements Expr {
    Expr a, b;
    PlusExpr(Expr a, Expr b) { this.a = a; this.b = b; }
    public int eval() { return a.eval() + b.eval(); }
}
final class TimesExpr implements Expr {
    Expr a, b;
    TimesExpr(Expr a, Expr b) { this.a = a; this.b = b; }
    public int eval() { return a.eval() * b.eval(); }
final class NegExpr implements Expr {
    Expr e;
   NegExpr(Expr e) { this.e = e; }
    public int eval() { return -e.eval(); }
}
```

Record Classes as Permitted Subclasses

You can name a record class in the permits clause of a sealed class or interface. See Record Classes for more information.

Record classes are implicitly final, so you can implement the previous example with record classes instead of ordinary classes:

package com.example.records.expressions;

```
public class TestExpressions {
  public static void main(String[] args) {
    // (6 + 7) * -8
    System.out.println(
      new TimesExpr(
        new PlusExpr(new ConstantExpr(6), new ConstantExpr(7)),
        new NegExpr(new ConstantExpr(8))
      ).eval());
   }
sealed interface Expr
    permits ConstantExpr, PlusExpr, TimesExpr, NegExpr {
    public int eval();
record ConstantExpr(int i) implements Expr {
    public int eval() { return i(); }
record PlusExpr(Expr a, Expr b) implements Expr {
    public int eval() { return a.eval() + b.eval(); }
record TimesExpr(Expr a, Expr b) implements Expr {
    public int eval() { return a.eval() * b.eval(); }
record NegExpr(Expr e) implements Expr {
    public int eval() { return -e.eval(); }
```

Narrowing Reference Conversion and Disjoint Types

Narrowing reference conversion is one of the conversions used in type checking cast expressions. It enables an expression of a reference type $\mathcal S$ to be treated as an expression of a different reference type $\mathcal T$, where $\mathcal S$ is not a subtype of $\mathcal T$. A narrowing reference conversion may require a test at run time to validate that a value of type $\mathcal S$ is a legitimate value of type $\mathcal T$. However, there are restrictions that prohibit conversion between certain pairs of types when it can be statically proven that no value can be of both types.

Consider the following example:

```
public interface Polygon { }
public class Rectangle implements Polygon { }

public void work(Rectangle r) {
    Polygon p = (Polygon) r;
}
```



The cast expression $Polygon\ p = (Polygon)\ r$ is permitted because it's possible that the Rectangle value r could be of type Polygon; Rectangle is a subtype of Polygon. However, consider this example:

```
public interface Polygon { }
public class Triangle { }

public void work(Triangle t) {
    Polygon p = (Polygon) t;
}
```

Even though the class Triangle and the interface Polygon are unrelated, the cast expression Polygon p = (Polygon) t is also permitted because at run time these types could be related. A developer could declare the following class:

```
class MeshElement extends Triangle implements Polygon { }
```

However, there are cases where the compiler can deduce that there are no values (other than the null reference) shared between two types; these types are considered *disjoint*. For example:

Because the class UtahTeapot is final, it's impossible for a class to be a descendant of both Polygon and UtahTeapot. Therefore, Polygon and UtahTeapot are disjoint, and the cast statement Polygon p = (Polygon) u isn't permitted.

The compiler has been enhanced to navigate any sealed hierarchy to check if your cast statements are permitted. For example:

```
public sealed interface Shape permits Polygon { }
public non-sealed interface Polygon extends Shape { }
public final class UtahTeapot { }
public class Ring { }

public void work(Shape s) {
    UtahTeapot u = (UtahTeapot) s; // Error
    Ring r = (Ring) s; // Permitted
}
```

The first cast statement u = (UtahTeapot) s isn't permitted; a Shape can only be a Polygon because Shape is sealed. However, as Polygon is non-sealed, it can be extended. However, no potential subtype of Polygon can extend utahTeapot as utahTeapot is final. Therefore, it's impossible for a Shape to be a utahTeapot.

In contrast, the second cast statement Ring r = (Ring) s is permitted; it's possible for a Shape to be a Ring because Ring is not a final class.

APIs Related to Sealed Classes and Interfaces

The class java.lang.Class has two new methods related to sealed classes and interfaces:

- java.lang.constant.ClassDesc[] permittedSubclasses(): Returns an array containing java.lang.constant.ClassDesc objects representing all the permitted subclasses of the class if it is sealed; returns an empty array if the class is not sealed
- boolean isSealed(): Returns true if the given class or interface is sealed



Pattern Matching for instanceof

Pattern matching involves testing whether an object has a particular structure, then extracting data from that object if there's a match. You can already do this with Java; however, pattern matching introduces new language enhancements that enable you to conditionally extract data from objects with code that's more concise and robust.

For background information about pattern matching for the instanceof operator, see JEP 394.

Consider the following code that calculates the perimeter of certain shapes:

```
public interface Shape {
    public static double getPerimeter(Shape shape) throws
IllegalArgumentException {
        if (shape instanceof Rectangle) {
            Rectangle r = (Rectangle) shape;
            return 2 * r.length() + 2 * r.width();
        } else if (shape instanceof Circle) {
            Circle c = (Circle) shape;
            return 2 * c.radius() * Math.PI;
        } else {
            throw new IllegalArgumentException("Unrecognized shape");
    }
public class Rectangle implements Shape {
    final double length;
    final double width;
   public Rectangle(double length, double width) {
        this.length = length;
        this.width = width;
   double length() { return length; }
   double width() { return width; }
public class Circle implements Shape {
    final double radius;
   public Circle(double radius) {
        this.radius = radius;
    double radius() { return radius; }
```

The method getPerimeter performs the following:

A test to determine the type of the Shape object

- 2. A conversion, casting the Shape object to Rectangle or Circle, depending on the result of the instanceof operator
- A destructuring, extracting either the length and width or the radius from the Shape object

Pattern matching enables you to remove the conversion step by changing the second operand of the <code>instanceof</code> operator with a type pattern, making your code shorter and easier to read:

```
public static double getPerimeter(Shape shape) throws
IllegalArgumentException {
    if (shape instanceof Rectangle r) {
        return 2 * r.length() + 2 * r.width();
    } else if (shape instanceof Circle c) {
        return 2 * c.radius() * Math.PI;
    } else {
        throw new IllegalArgumentException("Unrecognized shape");
    }
}
```

Note:

Removing this conversion step also makes your code safer. Testing an object's type with the <code>instanceof</code>, then assigning that object to a new variable with a cast can introduce coding errors in your application. You might change the type of one of the objects (either the tested object or the new variable) and accidentally forget to change the type of the other object.

A *pattern* is a combination of a predicate, or test, that can be applied to a target and a set of local variables, called *pattern variables*, that are assigned values extracted from the target only if the test is successful. The *predicate* is a Boolean-valued function of one argument; in this case, it's the instanceof operator testing whether the Shape argument is a Rectangle or a Circle. The *target* is the argument of the predicate, which is the Shape value. The pattern variables are those that store data from the target only if the predicate returns true, which are the variables r and s.

A *type pattern* consists of a predicate that specifies a type, along with a single pattern variable. In this example, the type patterns are Rectangle r and Circle c.

Scope of Pattern Variables

The scope of a pattern variable are the places where the program can reach only if the instanceof operator is true:

```
public static double getPerimeter(Shape shape) throws
IllegalArgumentException {
    if (shape instanceof Rectangle s) {
        // You can use the pattern variable s (of type Rectangle)
here.
} else if (shape instanceof Circle s) {
        // You can use the pattern variable s of type Circle here
        // but not the pattern variable s of type Rectangle.
} else {
```

```
// You cannot use either pattern variable here. \} \label{eq:cannot use} \}
```

The scope of a pattern variable can extend beyond the statement that introduced it:

```
public static boolean bigEnoughRect(Shape s) {
   if (!(s instanceof Rectangle r)) {
        // You cannot use the pattern variable r here because
        // the predicate s instanceof Rectangle is false.
        return false;
   }
   // You can use r here.
   return r.length() > 5;
}
```

You can use a pattern variable in the expression of an if statement:

```
if (shape instanceof Rectangle r && r.length() > 5) {
    // ...
}
```

Because the conditional-AND operator (&&) is short-circuiting, the program can reach the r.length() > 5 expression only if the instanceof operator is true.

Conversely, you can't pattern match with the instanceof operator in this situation:

```
if (shape instanceof Rectangle r || r.length() > 0) { // error // ...}
```

The program can reach the $r.length() \mid \mid$ 5 if the instanceof is false; thus, you cannot use the pattern variable r here.

Record Classes

Record classes, which are a special kind of class, help to model plain data aggregates with less ceremony than normal classes.

For background information about record classes, see JEP 395.

A record declaration specifies in a header a description of its contents; the appropriate accessors, constructor, equals, hashCode, and toString methods are created automatically. A record's fields are final because the class is intended to serve as a simple "data carrier".

For example, here is a record class with two fields:

```
record Rectangle(double length, double width) { }
```

This concise declaration of a rectangle is equivalent to the following normal class:

```
public final class Rectangle {
   private final double length;
   private final double width;
   public Rectangle(double length, double width) {
        this.length = length;
        this.width = width;
    }
   double length() { return this.length; }
   double width() { return this.width; }
    // Implementation of equals() and hashCode(), which specify
    // that two record objects are equal if they
    // are of the same type and contain equal field values.
   public boolean equals...
   public int hashCode...
    // An implementation of toString() that returns a string
    // representation of all the record class's fields,
    // including their names.
   public String toString() {...}
```

A record class declaration consists of a name; optional type parameters (generic record declarations are supported); a header, which lists the "components" of the record; and a body.

A record class declares the following members automatically:

• For each component in the header, the following two members:

- A private final field with the same name and declared type as the record component. This field is sometimes referred to as a *component field*.
- A public accessor method with the same name and type of the component; in the Rectangle record class example, these methods are Rectangle::length() and Rectangle::width().
- A canonical constructor whose signature is the same as the header. This
 constructor assigns each argument from the new expression that instantiates the
 record class to the corresponding component field.
- Implementations of the equals and hashCode methods, which specify that two
 record classes are equal if they are of the same type and contain equal
 component values.
- An implementation of the toString method that includes the string representation of all the record class's components, with their names.

As record classes are just special kinds of classes, you create a record object (an instance of a record class) with the new keyword, for example:

```
Rectangle r = new Rectangle(4,5);
```

The Canonical Constructor of a Record Class

The following example explicitly declares the canonical constructor for the Rectangle record class. It verifies that length and width are greater than zero. If not, it throws an IllegalArgumentException:

Repeating the record class's components in the signature of the canonical constructor can be tiresome and error-prone. To avoid this, you can declare a *compact constructor* whose signature is implicit (derived from the components automatically).

For example, the following compact constructor declaration validates <code>length</code> and <code>width</code> in the same way as in the previous example:



```
}
```

This succinct form of constructor declaration is only available in a record class. Note that the statements this.length = length; and this.width = width; which appear in the canonical constructor do not appear in the compact constructor. At the end of a compact constructor, its implicit formal parameters are assigned to the record class's private fields corresponding to its components.

Explicit Declaration of Record Class Members

You can explicitly declare any of the members derived from the header, such as the <code>public</code> accessor methods that correspond to the record class's components, for example:

```
record Rectangle(double length, double width) {
    // Public accessor method
    public double length() {
        System.out.println("Length is " + length);
        return length;
    }
}
```

If you implement your own accessor methods, then ensure that they have the same characteristics as implicitly derived accessors (for example, they're declared public and have the same return type as the corresponding record class component). Similarly, if you implement your own versions of the equals, hashCode, and toString methods, then ensure that they have the same characteristics and behavior as those in the java.lang.Record class, which is the common superclass of all record classes.

You can declare static fields, static initializers, and static methods in a record class, and they behave as they would in a normal class, for example:

```
record Rectangle(double length, double width) {
    // Static field
    static double goldenRatio;

    // Static initializer
    static {
        goldenRatio = (1 + Math.sqrt(5)) / 2;
    }

    // Static method
    public static Rectangle createGoldenRectangle(double width) {
        return new Rectangle(width, width * goldenRatio);
    }
}
```

You cannot declare instance variables (non-static fields) or instance initializers in a record class.



For example, the following record class declaration doesn't compile:

```
record Rectangle(double length, double width) {
    // Field declarations must be static:
    BiFunction<Double, Double, Double> diagonal;

    // Instance initializers are not allowed in records:
    {
        diagonal = (x, y) -> Math.sqrt(x*x + y*y);
    }
}
```

You can declare instance methods in a record class, independent of whether you implement your own accessor methods. You can also declare nested classes and interfaces in a record class, including nested record classes (which are implicitly static). For example:

You cannot declare native methods in a record class.

Features of Record Classes

A record class is implicitly final, so you cannot explicitly extend a record class. However, beyond these restrictions, record classes behave like normal classes:

You can create a generic record class, for example:

```
record Triangle<C extends Coordinate> (C top, C left, C right) { }
```

 You can declare a record class that implements one or more interfaces, for example:

```
record Customer(...) implements Billable { }
```



You can annotate a record class and its individual components, for example:

```
import java.lang.annotation.*;
@Retention(RetentionPolicy.RUNTIME)
@Target(ElementType.FIELD)
public @interface GreaterThanZero { }

record Rectangle(
    @GreaterThanZero double length,
    @GreaterThanZero double width) { }
```

If you annotate a record component, then the annotation may be propagated to members and constructors of the record class. This propagation is determined by the contexts in which the annotation interface is applicable. In the previous example, the <code>@Target(ElementType.FIELD)</code> meta-annotation means that the <code>@GreaterThanZero</code> annotation is propagated to the field corresponding to the record component. Consequently, this record class declaration would be equivalent to the following normal class declaration:

```
public final class Rectangle {
    private final @GreaterThanZero double length;
    private final @GreaterThanZero double width;

    public Rectangle(double length, double width) {
        this.length = length;
        this.width = width;
    }

    double length() { return this.length; }
    double width() { return this.width; }
}
```

Record Classes and Sealed Classes and Interfaces

Record classes work well with sealed classes and interfaces. See Record Classes as Permitted Subclasses for an example.

Local Record Classes

A local record class is similar to a local class; it's a record class defined in the body of a method.

In the following example, a merchant is modeled with a record class, Merchant. A sale made by a merchant is also modeled with a record class, Sale. Both Merchant and Sale are top-level record classes. The aggregation of a merchant and their total monthly sales is modeled with a *local* record class, MonthlySales, which is declared inside the findTopMerchants method. This local record class improves the readability of the stream operations that follow:

```
import java.time.*;
import java.util.*;
import java.util.stream.*;
```



```
record Merchant(String name) { }
record Sale(Merchant merchant, LocalDate date, double value) { }
public class MerchantExample {
    List<Merchant> findTopMerchants(
        List<Sale> sales, List<Merchant> merchants, int year, Month
month) {
        // Local record class
        record MerchantSales(Merchant merchant, double sales) {}
        return merchants.stream()
            .map(merchant -> new MerchantSales(
                merchant, this.computeSales(sales, merchant, year,
month)))
            .sorted((m1, m2) -> Double.compare(m2.sales(), m1.sales()))
            .map(MerchantSales::merchant)
            .collect(Collectors.toList());
    double computeSales(List<Sale> sales, Merchant mt, int yr, Month
mo) {
        return sales.stream()
            .filter(s -> s.merchant().name().equals(mt.name()) &&
                s.date().getYear() == yr &&
                s.date().getMonth() == mo)
            .mapToDouble(s -> s.value())
            .sum();
    }
    public static void main(String[] args) {
        Merchant sneha = new Merchant("Sneha");
        Merchant raj = new Merchant("Raj");
        Merchant florence = new Merchant("Florence");
        Merchant leo = new Merchant("Leo");
        List<Merchant> merchantList = List.of(sneha, raj, florence,
leo);
        List<Sale> salesList = List.of(
            new Sale(sneha,
                               LocalDate.of(2020, Month.NOVEMBER, 13),
11034.20),
                              LocalDate.of(2020, Month.NOVEMBER, 20),
            new Sale(raj,
8234.23),
           new Sale(florence, LocalDate.of(2020, Month.NOVEMBER, 19),
10003.67),
            // ...
            new Sale(leo,
                             LocalDate.of(2020, Month.NOVEMBER, 4),
9645.34));
        MerchantExample app = new MerchantExample();
```

Like nested record classes, local record classes are implicitly static, which means that their own methods can't access any variables of the enclosing method, unlike local classes, which are never static.

Static Members of Inner Classes

Prior to Java SE 16, you could not declare an explicitly or implicitly static member in an inner class unless that member is a constant variable. This means that an inner class cannot declare a record class member because nested record classes are implicitly static.

In Java SE 16 and later, an inner class may declare members that are either explicitly or implicitly static, which includes record class members. The following example demonstrates this:

```
public class ContactList {
   record Contact(String name, String number) { }
   public static void main(String[] args) {
        class Task implements Runnable {
            // Record class member, implicitly static,
            // declared in an inner class
            Contact c;
            public Task(Contact contact) {
                c = contact;
            public void run() {
                System.out.println(c.name + ", " + c.number);
            }
        List<Contact> contacts = List.of(
            new Contact("Sneha", "555-1234"),
            new Contact("Raj", "555-2345"));
        contacts.stream()
                .forEach(cont -> new Thread(new Task(cont)).start());
    }
```



Record Serialization

You can serialize and deserialize instances of record classes, but you can't customize the process by providing writeObject, readObject, readObjectNoData, writeExternal, or readExternal methods. The components of a record class govern serialization, while the canonical constructor of a record class governs deserialization. See Serializable Records for more information and an extended example. See also the section Serialization of Records in the Java Object Serialization Specification.

APIs Related to Record Classes

The abstract class java.lang.Record is the common superclass of all record classes.

You might get a compiler error if your source file imports a class named Record from a package other than <code>java.lang</code>. A Java source file automatically imports all the types in the <code>java.lang</code> package though an implicit <code>import java.lang.*;</code> statement. This includes the <code>java.lang.Record class</code>, regardless of whether preview features are enabled or disabled.

Consider the following class declaration of com.myapp.Record:

```
package com.myapp;

public class Record {
    public String greeting;
    public Record(String greeting) {
        this.greeting = greeting;
    }
}
```

The following example, org.example.MyappPackageExample, imports com.myapp.Record with a wildcard but doesn't compile:

```
package org.example;
import com.myapp.*;
public class MyappPackageExample {
    public static void main(String[] args) {
        Record r = new Record("Hello world!");
    }
}
```

The compiler generates an error message similar to the following:



```
ambiguous
```

```
Record r = new Record("Hello world!");
```

both class com.myapp.Record in com.myapp and class java.lang.Record in java.lang match

Both Record in the com.myapp package and Record in the java.lang package are imported with a wildcard. Consequently, neither class takes precedence, and the compiler generates an error when it encounters the use of the simple name Record.

To enable this example to compile, change the import statement so that it imports the fully qualified name of Record:

import com.myapp.Record;



The introduction of classes in the java.lang package is rare but necessary from time to time, such as Enum in Java SE 5, Module in Java SE 9, and Record in Java SE 14.

The class java.lang.Class has two methods related to record classes:

- RecordComponent[] getRecordComponents(): Returns an array of java.lang.reflect.RecordComponent objects, which correspond to the record class's components.
- boolean isRecord(): Similar to isEnum() except that it returns true if the class was declared as a record class.



Switch Expressions

Like all expressions, switch expressions evaluate to a single value and can be used in statements. They may contain "case $\, L \, \rightarrow$ " labels that eliminate the need for break statements to prevent fall through. You can use a yield statement to specify the value of a switch expression.

For background information about the design of switch expressions, see JEP 361.

"case L ->" Labels

Consider the following switch statement that prints the number of letters of a day of the week:

```
public enum Day { SUNDAY, MONDAY, TUESDAY,
   WEDNESDAY, THURSDAY, FRIDAY, SATURDAY; }
// ...
    int numLetters = 0;
   Day day = Day.WEDNESDAY;
   switch (day) {
        case MONDAY:
        case FRIDAY:
        case SUNDAY:
           numLetters = 6;
           break;
        case TUESDAY:
            numLetters = 7;
            break;
        case THURSDAY:
        case SATURDAY:
           numLetters = 8;
           break;
        case WEDNESDAY:
            numLetters = 9;
            break;
        default:
            throw new IllegalStateException("Invalid day: " + day);
    System.out.println(numLetters);
```

It would be better if you could "return" the length of the day's name instead of storing it in the variable numLetters; you can do this with a switch expression. Furthermore, it would be better if you didn't need break statements to prevent fall through; they are laborious to write and easy to forget. You can do this with a new kind of case label.

The following is a switch expression that uses the new kind of case label to print the number of letters of a day of the week:

```
Day day = Day.WEDNESDAY;
System.out.println(
    switch (day) {
        case MONDAY, FRIDAY, SUNDAY -> 6;
        case TUESDAY -> 7;
        case THURSDAY, SATURDAY -> 8;
        case WEDNESDAY -> 9;
        default -> throw new IllegalStateException("Invalid day: "
+ day);
    }
);
```

The new kind of case label has the following form:

```
case label_1, label_2, ..., label_n -> expression: |throw-statementi| block
```

When the Java runtime matches any of the labels to the left of the arrow, it runs the code to the right of the arrow and does not fall through; it does not run any other code in the switch expression (or statement). If the code to the right of the arrow is an expression, then the value of that expression is the value of the switch expression.

You can use the new kind of case label in switch statements. The following is like the first example, except it uses "case $L \rightarrow$ " labels instead of "case L:" labels:

```
int numLetters = 0;
Day day = Day.WEDNESDAY;
switch (day) {
    case MONDAY, FRIDAY, SUNDAY -> numLetters = 6;
    case TUESDAY -> numLetters = 7;
    case THURSDAY, SATURDAY -> numLetters = 8;
    case WEDNESDAY -> numLetters = 9;
    default -> throw new IllegalStateException("Invalid day: " + day);
};
System.out.println(numLetters);
```

A "case L ->" label along with its code to its right is called a switch labeled rule.

"case L:" Statements and the yield Statement

You can use "case L:" labels in switch expressions; a "case L:" label along with its code to the right is called a switch labeled statement group:

```
Day day = Day.WEDNESDAY;
int numLetters = switch (day) {
   case MONDAY:
   case FRIDAY:
   case SUNDAY:
       System.out.println(6);
```



```
yield 6;
case TUESDAY:
    System.out.println(7);
    yield 7;
case THURSDAY:
    case SATURDAY:
        System.out.println(8);
        yield 8;
case WEDNESDAY:
        System.out.println(9);
        yield 9;
    default:
        throw new IllegalStateException("Invalid day: " + day);
};
System.out.println(numLetters);
```

The previous example uses yield statements. They take one argument, which is the value that the case label produces in a switch expression.

The yield statement makes it easier for you to differentiate between switch statements and switch expressions. A switch statement, but not a switch expression, can be the target of a break statement. Conversely, a switch expression, but not a switch statement, can be the target of a yield statement.



Note:

It's recommended that you use "case $\, L \, -$ " labels. It's easy to forget to insert break or yield statements when using "case $\, L \colon$ " labels; if you do, you might introduce unintentional fall through in your code.

For "case L ->" labels, to specify multiple statements or code that are not expressions or throw statements, enclose them in a block. Specify the value that the case label produces with the yield statement:

```
int numLetters = switch (day) {
        case MONDAY, FRIDAY, SUNDAY -> {
            System.out.println(6);
            yield 6;
        case TUESDAY -> {
            System.out.println(7);
            yield 7;
        case THURSDAY, SATURDAY -> {
            System.out.println(8);
            yield 8;
        case WEDNESDAY -> {
            System.out.println(9);
            yield 9;
       default -> {
            throw new IllegalStateException("Invalid day: " +
day);
   };
```

Exhaustiveness

Unlike switch statements, the cases of switch expressions must be *exhaustive*, which means that for all possible values, there must be a matching switch label. Thus, switch expressions normally require a default clause. However, for enum switch expressions that cover all known constants, the compiler inserts an implicit default clause.

In addition, a switch expression must either complete normally with a value or complete abruptly by throwing an exception. For example, the following code doesn't compile because the switch labeled rule doesn't contain a yield statement:

```
int i = switch (day) {
   case MONDAY -> {
       System.out.println("Monday");
       // ERROR! Block doesn't contain a yield statement
   }
   default -> 1;
};
```



The following example doesn't compile because the switch labeled statement group doesn't contain a yield statement:

```
i = switch (day) {
    case MONDAY, TUESDAY, WEDNESDAY:
        yield 0;
    default:
        System.out.println("Second half of the week");
        // ERROR! Group doesn't contain a yield statement
};
```

Because a switch expression must evaluate to a single value (or throw an exception), you can't jump through a switch expression with a break, yield, return, or continue statement, like in the following example:



7

Text Blocks

See Programmer's Guide to Text Blocks for more information about this language feature. For background information about text blocks, see JEP 378.



Local Variable Type Inference

In JDK 10 and later, you can declare local variables with non-null initializers with the var identifier, which can help you write code that's easier to read.

Consider the following example, which seems redundant and is hard to read:

```
URL url = new URL("http://www.oracle.com/");
URLConnection conn = url.openConnection();
Reader reader = new BufferedReader(
    new InputStreamReader(conn.getInputStream()));
```

You can rewrite this example by declaring the local variables with the var identifier. The type of the variables are inferred from the context:

```
var url = new URL("http://www.oracle.com/");
var conn = url.openConnection();
var reader = new BufferedReader(
    new InputStreamReader(conn.getInputStream()));
```

var is a reserved type name, not a keyword, which means that existing code that uses var as a variable, method, or package name is not affected. However, code that uses var as a class or interface name is affected and the class or interface needs to be renamed.

var can be used for the following types of variables:

Local variable declarations with initializers:

Enhanced for-loop indexes:

```
List<String> myList = Arrays.asList("a", "b", "c");
for (var element : myList) {...} // infers String
```

• Index variables declared in traditional for loops:

```
for (var counter = 0; counter < 10; counter++) \{...\} // infers int
```



try-with-resources variable:

```
try (var input =
    new FileInputStream("validation.txt")) {...} // infers
FileInputStream
```

 Formal parameter declarations of implicitly typed lambda expressions: A lambda expression whose formal parameters have inferred types is implicitly typed:

```
BiFunction<Integer, Integer, Integer> = (a, b) -> a + b;
```

In JDK 11 and later, you can declare each formal parameter of an implicitly typed lambda expression with the var identifier:

```
(var a, var b) \rightarrow a + b;
```

As a result, the syntax of a formal parameter declaration in an implicitly typed lambda expression is consistent with the syntax of a local variable declaration; applying the var identifier to each formal parameter in an implicitly typed lambda expression has the same effect as not using var at all.

You cannot mix inferred formal parameters and var-declared formal parameters in implicitly typed lambda expressions nor can you mix var-declared formal parameters and manifest types in explicitly typed lambda expressions. The following examples are not permitted:

Local Variable Type Inference Style Guidelines

Local variable declarations can make code more readable by eliminating redundant information. However, it can also make code less readable by omitting useful information. Consequently, use this feature with judgment; no strict rule exists about when it should and shouldn't be used.

Local variable declarations don't exist in isolation; the surrounding code can affect or even overwhelm the effects of var declarations. Style Guidelines for Local Variable Type Inference in Java examines the impact that surrounding code has on var declarations, explains tradeoffs between explicit and implicit type declarations, and provides guidelines for the effective use of var declarations.

