



8

Inner Classes

CERTIFICATION OBJECTIVES

- Inner Classes
- Method-Local Inner Classes
- Anonymous Inner Classes
- Static Nested Classes
- ✓ Two-Minute Drill

Q&A Self Test

Inner classes (including static nested classes) appear throughout the exam. Although there are no official exam objectives specifically about inner classes, the objectives related to declarations (1.2 and 4.1) and instantiation (6.3) include inner classes. More importantly, the code used to represent questions on virtually *any* topic on the exam can involve inner classes. Unless you deeply understand the rules and syntax for inner classes, you're likely to miss questions you'd otherwise be able to answer. *As if the exam weren't already tough enough.*

This chapter looks at the ins and outs (inners and outers?) of inner classes, and exposes you to the kinds of (often strange-looking) syntax examples you'll see scattered throughout the entire exam. So you've really got two goals for this chapter—to learn what you'll need to answer questions testing your inner class knowledge, and to learn how to read and understand inner class code so that you can correctly process questions testing your knowledge of *other* topics.

So what's all the hoopla about inner classes? Before we get into it, we have to warn you (if you don't already know) that inner classes have inspired passionate love 'em or hate 'em debates since first introduced in version 1.1 of the language. For once, we're going to try to keep our opinions to ourselves here and just present the facts as you'll need to know them for the exam. It's up to you to decide how—and to what extent—you should use them in your own development. We mean it. Not even our tone will betray our true feelings about them. (OK, OK, we'll tell you! We believe they have some powerful, efficient uses in very specific situations, including code that's easier to read and maintain, but they can also be abused and lead to code that's as clear as a cornfield maze, and to the syndrome known as “reuseless”...*code that's useless over and over again.*)

Inner classes let you define one class within another. They provide a type of scoping for your classes since you can make one class *a member of another class*. Just as classes have member *variables* and *methods*, a class can also have member *classes*. They come in several flavors, depending on how and where you define the inner class, including a special kind of inner class known as a “top-level nested class” (an inner class marked `static`), which technically isn't really an inner class. Because a static nested class is still a class defined within the scope of another class, we're still going to cover them in this chapter on inner classes.

Unlike the other chapters in this book, the certification objectives for inner classes don't have official exam objective numbers since they're part of other objectives covered elsewhere. So for this chapter, the Certification Objective headings represent the four inner class *topics* discussed in this chapter, rather than four official exam *objectives*.

- Inner Classes
- Method-local Inner Classes
- Anonymous Inner Classes
- Static Nested Classes

CERTIFICATION OBJECTIVE

Inner Classes

You're an OO programmer, so you know that for reuse and flexibility/extensibility you need to keep your classes specialized. In other words, a class should have code *only* for the things an object of that particular type needs to do; any *other* behavior should be part of another class better suited for *that* job. Sometimes, though, you find yourself designing a class where you discover you need behavior that belongs in a separate, specialized class, but also needs to be intimately tied to the class you're designing.

Event handlers are perhaps the best example of this (and in fact, one of the main reasons inner classes were added to the language in the first place). If you have a GUI class that performs some job like, say, a chat client, you might want the chat-client-specific methods (accept input, read new messages from server, send user input back to server, etc.) to be in the class. But how do those methods get invoked in the first place? A user clicks a button. Or types some text in the input field. Or a separate thread doing the I/O work of getting messages from the server has messages that need to be displayed in the GUI. So you have chat-client-specific methods, but you also need methods for handling the “events” (button presses, keyboard typing, I/O available, etc.) that drive the calls on those chat-client methods. The ideal scenario—from an OO perspective—is to keep the chat-client-specific methods in the ChatClient class, and put the event-handling *code* in a separate event-handling *class*.

Nothing unusual about that so far; after all, that's how you're *supposed* to design OO classes. As *specialists*. But here's the problem with the chat-client scenario: the event-handling code is intimately tied to the chat-client-specific code! Think about it: when the user presses a Send button (indicating that they want their typed-in message to be sent to the chat server), the chat-client code that sends the message

needs to read from a *particular* text field. In other words, if the user clicks Button A, the program is supposed to extract the text from the TextField B, *of a particular ChatClient instance*. Not from some *other* text field from some *other* object, but specifically the text field that a specific instance of the ChatClient class has a reference to. So the event-handling code needs access to the members of the ChatClient object, in order to be useful as a “helper” to a particular ChatClient instance.

And what if the ChatClient class needs to inherit from one class, but the event-handling code is better off inheriting from some *other* class? You can’t make a class extend from more than once class, so putting all the code (the chat-client-specific code and the event-handling code) in one class won’t work in that case. So what you’d really like to have is the benefit of putting your event code in a separate class (better OO, encapsulation, and the ability to extend a class other than the class the ChatClient extends) but yet still allow the event-handling code to have easy access to the members of the ChatClient (so the event-handling code can, for example, update the ChatClient’s private instance variables). You *could* manage it by making the members of the ChatClient accessible to the event-handling class by, for example, marking them public. But that’s not a good solution either.

You already know where this is going—one of the key benefits of an inner class is the “special relationship” an *inner class instance* shares with *an instance of the outer class*. That “special relationship” gives code in the inner class access to members of the enclosing (outer) class, *as if the inner class were part of the outer class*. In fact, that’s exactly what it means: the inner class *is* a part of the outer class. Not just a “part” but a full-fledged, card-carrying *member* of the outer class. Yes, an inner class instance has access to all members of the outer class, even *those marked private*. (Relax, that’s the whole point, remember? We want this separate inner class instance to have an intimate relationship with the outer class instance, but we still want to keep everyone *else* out. And besides, if you wrote the outer class, then you also wrote the inner class! So you’re not violating encapsulation; you *designed* it this way.)

Coding a “Regular” Inner Class

We use the term *regular* here to represent inner classes that are *not*

- static
- method-local
- anonymous

For the rest of this section, though, we'll just use the term *inner class* and drop the *regular*. (When we switch to one of the other three types in the preceding list, you'll know it.) You define an inner class within the curly braces of the outer class, as follows:

```
class MyOuter {
    class MyInner { }
}
```

Piece of cake. And if you compile it,

```
%javac MyOuter.java
```

you'll end up with *two* class files:

```
MyOuter.class
MyOuter$MyInner.class
```

The inner class is still, in the end, a separate class, so a class file is generated. But the inner class file isn't accessible to you in the usual way. You can't, for example, say

```
%java MyOuter$MyInner
```

in hopes of running the main method of the inner class, because a *regular* inner class can't have static declarations of any kind. *The only way you can access the inner class is through a live instance of the outer class!* In other words, only at runtime when there's already an instance of the outer class to tie the inner class instance to. You'll see all this in a moment. First, let's beef up the classes a little:

```
class MyOuter {
    private int x = 7;

    // inner class definition
    class MyInner {
        public void seeOuter() {
            System.out.println("Outer x is " + x);
        }
    } // close inner class definition

} // close outer class
```

The preceding code is perfectly legal. Notice that the inner class is indeed accessing a private member of the outer class. That's fine, because the inner class is also a member of the outer class. So just as any member of the outer class

(say, an instance method) can access any other member of the outer class, private or not, the inner class—also a member—can do the same.

OK, so now that we know how to write the code giving an inner class access to members of the outer class, how do you actually use it?

Instantiating an Inner Class

To instantiate an instance of an inner class, *you must have an instance of the outer class* to tie to the inner class. There are no exceptions to this rule; an inner class instance can never stand alone without a direct relationship with a specific instance of the outer class.

Instantiating an Inner Class from Within Code in the Outer Class

Most often, it is the outer class that creates instances of the inner class, since it is usually the outer class wanting to use the inner instance as a helper object, for its own personal, private use. We'll modify the `MyOuter` class to instantiate an instance of `MyInner`:

```
class MyOuter {
    private int x = 7;
    public void makeInner() {
        MyInner in = new MyInner();
        in.seeOuter();
    }

    class MyInner {
        public void seeOuter() {
            System.out.println("Outer x is " + x);
        }
    }
}
```

You can see in the preceding code that the `MyOuter` code treats `MyInner` just as though `MyInner` were any other accessible class—it instantiates it using the class name (`new MyInner()`), and then invokes a method on the reference variable (`in.seeOuter()`). But the only reason this syntax works is because the outer class instance method code is doing the instantiating. In other words, *there's already*

an instance of the outer class—the instance running the `makeInner()` method. So how do you instantiate a `MyInner` object from somewhere outside the `MyOuter` class? Is it even possible? (Well, since we're going to all the trouble of making a whole new subhead for it, as you'll see next, there's no big mystery here.)

Creating an Inner Class Object from Outside the Outer Class

Instance Code Whew. Long subhead there, but it does explain what we're trying to do. If we want to create an instance of the inner class, we must have an instance of the outer class. You already know that, but think about the implications... it means you can't instantiate the inner class from a `static` method of the outer class (because, don't forget, in static code *there is no this reference*) or from any other code in any other class. Inner class instances are always handed an implicit reference to the outer class. The compiler takes care of it, so you'll never see anything but the end result—the ability of the inner class to access members of the outer class. The code to make an instance from anywhere outside nonstatic code of the outer class is simple, but you must memorize this for the exam!

```
public static void main (String[] args) {
    MyOuter mo = new MyOuter();
    MyOuter.MyInner inner = mo.new MyInner();
    inner.seeOuter();
}
```

The preceding code is the same regardless of whether the `main()` method is within the `MyOuter` class or some *other* class (assuming the other class has access to `MyOuter`, and since `MyOuter` has default access, that means the code must be in a class within the same package as `MyOuter`).

If you're into one-liners, you can do it like this:

```
public static void main (String[] args) {
    MyOuter.MyInner inner = new MyOuter().new MyInner();
    inner.seeOuter();
}
```

You can think of this as though you're invoking a method on the outer instance, but the method happens to be a special inner class instantiation method, and it's invoked using the keyword `new`. Instantiating an inner class is the *only* scenario in which you'll invoke `new` *on* an instance as opposed to invoking `new` to *construct* an instance.

Here's a quick summary of the differences between inner class instantiation code that's *within* the outer class (but not static), and inner class instantiation code that's *outside* the outer class:

- **From inside the outer class instance code**, use the inner class name in the normal way:

```
MyInner mi = new MyInner();
```

- **From outside the outer class instance code (including static method code within the outer class)**, the inner class name must now include the outer class name,

```
MyOuter.MyInner
```

and to instantiate, you must use a reference to the outer class,

```
new MyOuter().new MyInner(); or outerObjRef.new MyInner();
```

if you already have an instance of the outer class.

Referencing the Inner or Outer Instance from Within the Inner Class

How does an object refer to itself normally? Using the `this` reference. Quick review of `this`:

- The keyword `this` can be used only from within instance code. In other words, not within static code.
- The `this` reference is a reference to the currently-executing object. In other words, the object whose reference was used to invoke the currently-running method.
- The `this` reference is the way an object can pass a reference to itself to some other code, as a method argument:

```
public void myMethod() {
    MyClass mc = new MyClass();
    mc.doStuff(this); // pass a ref to object running myMethod
}
```


Within an inner class code, the `this` reference refers to the instance of the inner class, as you'd probably expect, since `this` always refers to the currently-executing object. But what if the inner class code wants an explicit reference to the outer class instance the inner instance is tied to? In other words, *how do you reference the "outer this"?* Although normally the inner class code doesn't need a reference to the outer class, since it already has an implicit one it's using to access the members of the outer class, it would need a reference to the outer class if it needed to pass that reference to some other code as follows:

```
class MyInner {
    public void seeOuter() {
        System.out.println("Outer x is " + x);
        System.out.println("Inner class ref is " + this);
        System.out.println("Outer class ref is " + MyOuter.this);
    }
}
```

If we run the complete code as follows:

```
class MyOuter {
    private int x = 7;
    public void makeInner() {
        MyInner in = new MyInner();
        in.seeOuter();
    }
    class MyInner {
        public void seeOuter() {
            System.out.println("Outer x is " + x);
            System.out.println("Inner class ref is " + this);
            System.out.println("Outer class ref is " + MyOuter.this);
        }
    }
    public static void main (String[] args) {
        MyOuter.MyInner inner = new MyOuter().new MyInner();
        inner.seeOuter();
    }
}
```

the output is

```
Outer x is 7
Inner class ref is MyOuter$MyInner@113708
Outer class ref is MyOuter@33f1d7
```

So the rules for an inner class referencing itself or the outer instance are as follows:

- To reference the inner class instance itself, from *within* the inner class code, use `this`.
- To reference the “*outer this*” (the outer class instance) from within the inner class code, use `<NameOfOuterClass>.this` (example, `MyOuter.this`).

Member Modifiers Applied to Inner Classes A regular inner class is a member of the outer class just as instance variables and methods are, so the following modifiers can be applied to an inner class:

- `final`
- `abstract`
- `public`
- `private`
- `protected`
- `static`—*except `static` turns it into a top-level nested class rather than an inner class.*
- `strictfp`

CERTIFICATION OBJECTIVE

Method-Local Inner Classes

A regular inner class is scoped inside another class’ curly braces, but outside any method code (in other words, at the same level as an instance variable is declared). But you can also define an inner class within a method:

```
class MyOuter2 {
    private String x = "Outer2";

    void doStuff() {
        class MyInner {
            public void seeOuter() {
                System.out.println("Outer x is " + x);
            }
        }
    }
}
```

```

        } // close inner class method
    } // close inner class definition
} // close outer class method doStuff()

} // close outer class

```

The preceding code declares a class, `MyOuter2`, with one method, `doStuff()`. But *inside* `doStuff()`, another class, `MyInner`, is declared, and it has a method of its own, `seeOuter()`. The code above is completely useless, however, because *it never instantiates the inner class!* Just because you *declared* the class doesn't mean you created an *instance* of it. So if you want to actually *use* the inner class (say, to invoke its methods), then you must make an instance of it somewhere *within the method but below the inner class definition*. The following legal code shows how to instantiate and use a method-local inner class:

```

class MyOuter2 {
    private String x = "Outer2";
    void doStuff() {
        class MyInner {
            public void seeOuter() {
                System.out.println("Outer x is " + x);
            } // close inner class method
        } // close inner class definition
        MyInner mi = new MyInner(); // This line must come
                                    //after the class

        mi.seeOuter();
    } // close outer class method doStuff()
} // close outer class

```

What a Method-Local Inner Object Can and Can't Do

A method-local inner class can be instantiated only within the method where the inner class is defined. In other words, no other code running in any other method—inside or outside the outer class—can ever instantiate the method-local inner class. Like regular inner class objects, the method-local inner class object shares a special relationship with the enclosing (outer) class object, and can access its private (or any other) members. However, *the inner class object cannot use the local variables of the method the inner class is in.* Why not?

Think about it. The local variables of the method live on the stack, and exist only for the lifetime of the method. You already know that the scope of a local variable is limited to the method the variable is declared in. When the method ends, the stack

frame is blown away and the variable is history. But even after the method completes, the inner class object created within it might still be alive on the heap if, for example, a reference to it was passed into some other code and then stored in an instance variable. Because the local variables aren't guaranteed to be alive as long as the method-local inner class object, the inner class object can't use them. *Unless the local variables are marked final!* The following code attempts to access a local variable from within a method-local inner class:

```
class MyOuter2 {
    private String x = "Outer2";
    void doStuff() {
        String z = "local variable";
        class MyInner {
            public void seeOuter() {
                System.out.println("Outer x is " + x);
                System.out.println("Local variable z is " + z); // Won't Compile!
            } // close inner class method
        } // close inner class definition
    } // close outer class method doStuff()
} // close outer class
```

Compiling the preceding code *really* upsets the compiler:

```
MyOuter2.java:8: local variable z is accessed from within inner class;
needs to be declared final
        System.out.println("Local variable z is " + z);
                               ^
```

Marking the local variable *z* as `final` fixes the problem:

```
final String z = "local variable"; // Now inner object can use it
```

And just a reminder about modifiers within a method: the same rules apply to method-local inner classes as to local variable declarations. You can't, for example, mark a method-local inner class `public`, `private`, `protected`, `static`, `transient`, and the like. The only modifiers you *can* apply to a method-local inner class are `abstract` and `final`. (But of course, never both of those at the same time as with any other class or method.)

exam Watch

Remember that a local class declared in a static method has access to only static members of the enclosing class, since there is no associated instance of the enclosing class. If you're in a static method there is no 'this', so an inner class in a static method is subject to the same restrictions as the static method. In other words, no access to 'instance variables'.

CERTIFICATION OBJECTIVE

Anonymous Inner Classes

So far we've looked at defining a class within an enclosing class (a regular inner class), and within a method (a method-local inner class). Finally, we're going to look at the most unusual syntax you might ever see in Java, inner classes declared without any class name at all (hence the word *anonymous*). And if that's not weird enough, you can even define these classes not just within a method, but within an *argument* to a method. We'll look first at the plain old (as if there is such a thing as a *plain old* anonymous inner class) version (actually, even the plain old version comes in two flavors), and then at the argument-declared anonymous inner class.

Perhaps your most important job here is to *learn to not be thrown when you see the syntax*. The exam is littered with anonymous inner class code; you might see it on questions about threads, wrappers, overriding, garbage collection, and you get the idea.

Plain Old Anonymous Inner Classes, Flavor One

Check out the following legal-but-strange-the-first-time-you-see-it code:

```
class Popcorn {
    public void pop() {
        System.out.println("popcorn");
    }
}
class Food {
    Popcorn p = new Popcorn() {
        public void pop() {
            System.out.println("anonymous popcorn");
        }
    };
}
```

Let's look at what's in the preceding code:

- We define two classes, `Popcorn` and `Food`.
- `Popcorn` has one method, `pop()`.
- `Food` has one instance variable, declared as type `Popcorn`.
- That's it for `Food`. `Food` has *no* methods.

And here's the big thing to get:

- The `Popcorn` reference variable refers *not* to an instance of `Popcorn`, but to *an instance of an anonymous (unnamed) subclass of `Popcorn`*.

Let's look at just the anonymous class code:

```
2. Popcorn p = new Popcorn() {
3.     public void pop() {
4.         System.out.println("anonymous popcorn");
5.     }
6. };
```

Line 2 Line 2 starts out as an instance variable declaration of type `Popcorn`. But instead of looking like this:

```
Popcorn p = new Popcorn(); // ← notice the semicolon at the end
```

there's a curly brace at the end of line 2, where a semicolon would normally be

```
Popcorn p = new Popcorn() { // ← a curly brace rather than semicolon
```

You can read line 2 as saying,

“Declare a reference variable, *p*, of type `Popcorn`. Then declare a new class which has no name, but which is a *subclass* of `Popcorn`. And here's the curly brace that opens the class definition...”

Line 3 Line 3, then, is actually the first statement within the new class definition. And what is it doing? Overriding the `pop()` method of the superclass `Popcorn`. This is the whole point of making an anonymous inner class—to *override one or more methods of the superclass!* (Or to implement methods of an interface, but we'll save that for a little later.)

Line 4 Line 4 is the first (and in this case *only*) statement within the overriding `pop()` method. Nothing special there.

Line 5 Line 5 is the closing curly brace of the `pop()` method. Again, nothing special there.

Line 6 Here's where you have to pay attention: line 6 includes a *curly brace closing off the anonymous class definition* (it's the companion brace to the one on line 2), but there's more! Line 6 also has *the semicolon that ends the statement started on line 2*. The statement where it all began. The statement declaring and initializing the `Popcorn` reference variable. And what you're left with is a `Popcorn` reference to a brand new *instance* of a brand new, just-in-time, anonymous (no name) *subclass* of `Popcorn`.

exam
Watch

The closing semicolon is often hard to spot. So you might see code on the exam like this:

```
2. Popcorn p = new Popcorn() {
3.     public void pop() {
4.         System.out.println("anonymous popcorn");
5.     }
6. } // ← Missing the semicolon needed to end statement on 2!!
7. Foo f = new Foo();
```

You'll need to be especially careful about the syntax when inner classes are involved, because the code on Line 6 looks perfectly natural. We're not used to seeing semicolons following curly braces (the only other time it happens is with shortcut array initializations).

exam
Watch

Polymorphism is in play when anonymous inner classes are involved. Remember that, as in the preceding Popcorn example, we're using a superclass reference variable type to refer to a subclass object. What are the implications? You can only call methods on an anonymous inner class reference that are defined in the reference variable type! This is no different from any other polymorphic references, for example,

```
class Horse extends Animal{
    void buck() { }
}
class Animal {
    void eat() { }
}
class Test {
    public static void main (String[] args) {
        Animal h = new Horse();
        h.eat(); // Legal, class Animal has an eat() method
        h.buck(); // Not legal! Class Animal doesn't have buck()
    }
}
```

So on the exam, you must be able to spot an anonymous inner class that—rather than overriding a method of the superclass—defines its own new method. The method definition isn't the problem, though; the real issue is how do you invoke that new method? The reference variable type (the superclass) won't know anything about that new method (defined in the anonymous subclass), so the compiler will complain if you try to invoke any method on an anonymous inner class reference that is not in the superclass class definition.

Check out the following, illegal code:

```
class Popcorn {
    public void pop() {
        System.out.println("popcorn");
    }
}
class Food {
    Popcorn p = new Popcorn() {
        public void sizzle() {
            System.out.println("anonymous sizzling popcorn");
        }
        public void pop() {
```



```

        System.out.println("anonymous popcorn");
    }
};
public void popIt() {
    p.pop(); // OK, Popcorn has a pop() method
    p.sizzle(); // Not Legal! Popcorn does not have sizzle()
}
}

```

Compiling the preceding code gives us,

```

Anon.java:19: cannot resolve symbol
symbol  : method sizzle ()
location: class Popcorn
    p.sizzle();
      ^

```

which is the compiler's way of saying, "I can't find method `sizzle()` in class `Popcorn`," followed by, "Get a clue."

Plain Old Anonymous Inner Classes, Flavor Two

The only difference between flavor one and flavor two is that flavor one creates an anonymous *subclass* of the specified *class* type, whereas flavor two creates an anonymous *implementer* of the specified *interface* type. In the previous examples, we defined a new anonymous subclass of type `Popcorn` as follows:

```
Popcorn p = new Popcorn() {
```

But if `Popcorn` were an *interface* type instead of a *class* type, then the new anonymous class would be an *implementer* of the *interface* rather than a *subclass* of the *class*. Look at the following example:

```

interface Cookable {
    public void cook();
}
class Food {
    Cookable c = new Cookable() {
        public void cook() {
            System.out.println("anonymous cookable implementer");
        }
    };
}

```

The preceding code, like the Popcorn example, still creates an instance of an anonymous inner class, but this time the new just-in-time class is an implementer of the `Cookable` interface. And note that this is the only time you will ever see the syntax,

```
new Cookable()
```

where `Cookable` is an *interface* rather than a non-abstract class type. Because think about it, *you can't instantiate an interface*, yet that's what the code *looks* like it's doing. But of course it's not instantiating a `Cookable` object, it's creating an instance of a new, anonymous, implementer of `Cookable`. So you can read this line,

```
Cookable c = new Cookable() {
```

as “Declare a reference variable of type `Cookable` that, obviously, will refer to an object from a class that implements the `Cookable` interface. But, oh yes, we don't yet *have* a class that implements `Cookable`, so we're going to make one right here, right now. We don't need a name for the class, but it will be a class that implements `Cookable`, and this curly brace starts the definition of the new implementing class.”

One more thing to keep in mind about anonymous interface implementers—*they can implement only one interface*. There simply isn't any mechanism to say that your anonymous inner class is going to implement multiple interfaces. In fact, an anonymous inner class can't even extend a class and implement an interface at the same time. The inner class has to choose either to be a subclass of a named class—and not directly implement any interfaces at all—*or* to implement a single interface. By directly, we mean actually using the keyword `implements` as part of the class declaration. If the anonymous inner class is a subclass of a class type, it automatically becomes an implementer of any interfaces implemented by the superclass.

exam
Watch

Don't be fooled by any attempts to instantiate an interface except in the case of an anonymous inner class. The following is not legal,

```
Runnable r = new Runnable(); // can't instantiate interface
```

whereas the following is legal, because it's instantiating an implementer of the `Runnable` interface (an anonymous implementation class):

```
Runnable r = new Runnable() { // curly brace instead of semicolon
    public void run() { }
};
```

Argument-Defined Anonymous Inner Class

If you understood what we've covered so far in this chapter, then this last part will be simple. If you *are* still a little fuzzy on anonymous classes, however, then you should reread the previous sections. If they're not completely clear, we'd like to take full responsibility for the confusion. But we'll be happy to share.

OK, if you've made it to this sentence then we're all going to assume you understood the previous section, and now we're just going to add one new twist. Imagine the following scenario. You're typing along, creating the Perfect Class, when you write code calling a method on a `Bar` object, and that method takes an object of type `Foo` (an interface).

```
class MyWonderfulClass {
    void go() {
        Bar b = new Bar();
        b.doStuff(AckWeDon'tHaveAFoo!); // Don't try to compile this at home
    }
}
interface Foo {
    void foof();
}
class Bar {
    void doStuff(Foo f) { }
}
```

No *problem*, except that you don't *have* an object from a class that implements `Foo`. But you can't instantiate one, either, because *you don't even have a class that implements Foo*, let alone an instance of one. So you first need a class that implements `Foo`, and then you need an instance of that class to pass to the `Bar` class' `doStuff()` method. Savvy Java programmer that you are, you simply define an anonymous inner class, *right inside the argument*. That's right, just where you least expect to find a class. And here's what it looks like:

```
1. class MyWonderfulClass {
2.     void go() {
3.         Bar b = new Bar();
4.         b.doStuff(new Foo() {
5.             public void foof() {
6.                 System.out.println("foofy");
7.             } // end foof method
8.         }); // end inner class def, arg, and end statement
9.     } // end go()
10. } // end class
11.
12. interface Foo {
```

```

13.     void foof();
14. }
15. class Bar {
16.     void doStuff(Foo f) { }
17. }

```

All the action starts on line 4. We're calling `doStuff()` on a `Bar` object, but the method takes an instance that IS-A `Foo`, where `Foo` is an interface. So we must make both an *implementation* class and an *instance* of that class, all right here in the argument to `doStuff()`. So that's what we do. We write

```
new Foo() {
```

to start the new class definition for the anonymous class that implements the `Foo` interface. `Foo` has a single method to implement, `foof()`, so on lines 5, 6, and 7 we implement the `foof()` method. Then on line 8—whoa!—more strange syntax appears. The first curly brace closes off the new anonymous class definition. But don't forget that this all happened as part of a method argument, so the close parenthesis ')' finishes off the method invocation, and then we must still end the statement that began on line 4, so we end with a semicolon. Study this syntax! You *will* see anonymous inner classes on the exam, and you'll have to be very, very picky about the way they're closed. If they're argument local, they end like this,

```
});
```

but if they're just plain old anonymous classes, then they end like this:

```
};
```

Regardless, the syntax is not what you use in virtually any other part of Java, so be careful. Any question from any part of the exam might involve anonymous inner classes as part of the code.

CERTIFICATION OBJECTIVE

Static Nested Classes

We saved the easiest for last, as a kind of treat :)

You'll sometimes hear static nested classes referred to as top-level nested classes, or *static inner classes*, but they really aren't inner classes at all, by the standard definition

of an inner class. While an inner class (regardless of the flavor) enjoys that *special relationship* with the outer class (or rather the *instances* of the two classes share a relationship), a static nested class does not. It is simply a non-inner (also called “top-level”) class scoped within another. So with static classes it’s really more about name-space resolution than about an implicit relationship between the two classes.

A static nested class is simply *a class that’s a static member of the enclosing class*, as follows:

```
class BigOuter {
    static class Nested { }
}
```

The class itself isn’t really “static;” there’s no such thing as a static class. The `static` modifier in this case says that the nested class is *a static member of the outer class*. That means it can be accessed, as with other static members, *without having an instance of the outer class*.



Just as a static method does not have access to the instance variables and methods of the class, a static nested class does not have access to the instance variables and methods of the outer class. Look for static nested classes with code that behaves like a nonstatic (regular inner) class.

Instantiating a Static Nested Class

The syntax for instantiating a static nested class is a little different from a normal inner class, and looks like this:

```
class BigOuter {
    static class Nested { }
}
class Broom {
    public static void main (String [] args) {
        BigOuter.Nested n = new BigOuter.Nested(); //Use both class names
    }
}
```

CERTIFICATION SUMMARY

You’re on the home stretch now; just one more chapter follows this one. You’ve learned all about inner classes (including static nested classes), and you’re aware that they’ll show up throughout the exam, regardless of the topic. You’re comfortable with

the sometimes bizarre syntax, and you know how to spot legal and illegal inner class definitions.

We looked first at “regular” inner classes, where one class is a member of another. You learned that coding an inner class means putting the class definition of the inner class inside the curly braces of the enclosing (outer) class, but outside of any method or other code block. We learned that an inner class *instance* shares a special relationship with a specific *instance* of the outer class, and that this special relationship lets the inner class access all members of the outer class, including those marked `private`. You learned that to instantiate an inner class, you *must* have a reference to an instance of the outer class.

Next we looked at method-local inner classes—classes defined *inside* a method. We saw that the code for a method-local inner class looks virtually the same as the code for any other class definition, except that you can’t apply an access modifier the way you can to a regular inner class. You also learned why method-local inner classes cannot use nonfinal local variables declared within the same method—the inner class instance may outlive the stack frame, so the local variable might vanish while the inner class object is still alive. We showed you that to *use* the inner class you need to instantiate it, and that the instantiation must come *after* the class declaration in the method.

We also explored the strangest inner class type of all—the *anonymous* inner class. You learned that they come in two forms, normal and argument-local. Normal, ho-hum, anonymous inner classes are created as part of a variable assignment, while argument-local inner classes are actually declared, defined, and automatically instantiated *all within the argument to a method!* We covered the way anonymous inner classes can be either a *subclass* of the named class type, or an *implementer* of the named interface. Finally, we looked at how polymorphism applies to anonymous inner classes: you can invoke on the new instance only those methods defined in the named class or interface type. In other words, even if the anonymous inner class defines its own new method, no code from anywhere outside the inner class will be able to invoke that method.

As if we weren't already having enough fun for one day, we pushed on to static inner classes, which really aren't inner classes at all. Known as static nested classes (or top-level nested classes), a nested class marked with the `static` modifier is quite similar to any other non-inner class, except that to access it, code must have access to both the nested and enclosing class. We saw that because the class is static, no instance of the enclosing class is needed, and thus the static nested class *does not share a special relationship with any instance of the enclosing class*.

We've finished the inner class tour, and now it's up to you to review the Two-Minute Drill and *take the Self Test*. We can virtually guarantee that if you can't answer these questions correctly, you probably can't pass the exam. On the bright side, though, a strong knowledge of inner class syntax and behavior should get you through some of the exam's *toughest* questions. And did you know that recent studies show that intense mental effort—like working on logic puzzles—can actually increase the synaptic connections in your brain? *This* chapter's Self-Test qualifies as “intense mental effort,” and might be covered under some health insurance plans.

✓ TWO-MINUTE DRILL

Here are some of the key points from the certification objectives (which for this chapter, means the four inner class topics).

Inner Classes

- ❑ A “regular” inner class is declared *inside* the curly braces of another class, but *outside* any method or other code block.
- ❑ An inner class is a full-fledged member of the enclosing (outer) class, so it can be marked with an access modifier as well as the `abstract` or `final` modifiers (but of course, never both `abstract` and `final` together—remember that `abstract` means it *must* be subclassed, whereas `final` means it *cannot* be subclassed).
- ❑ An inner class instance shares a special relationship with an instance of the enclosing class. This relationship gives the inner class access to *all* of the outer class’ members, including those marked `private`.
- ❑ To instantiate an inner class, you must have a reference to an instance of the outer class.
- ❑ From code within the enclosing class, you can instantiate the inner class using only the name of the inner class, as follows:

```
MyInner mi = new MyInner ();
```
- ❑ From code outside the enclosing class’ instance methods, you can instantiate the inner class only by using both the inner and outer class names, and a reference to the outer class as follows:

```
MyOuter mo = new MyOuter ();  
MyOuter.MyInner inner = mo.new MyInner ();
```
- ❑ From code within the inner class, the keyword `this` holds a reference to the inner class instance. To reference the *outer this* (in other words, the instance of the outer class that this inner instance is tied to) precede the keyword `this` with the outer class name as follows:

```
MyOuter.this;
```


Method-Local Inner Classes

- ❑ A method-local inner class is defined within a method of the enclosing class.
- ❑ For the inner class to be used, you must instantiate it, and that instantiation must happen within the same method, but *after* the class definition code.
- ❑ A method-local inner class cannot use variables declared within the method (including parameters) unless those variables are marked `final`.
- ❑ The only modifiers you can apply to a method-local inner class are `abstract` and `final`. (Never both at the same time, though.)

Anonymous Inner Classes

- ❑ Anonymous inner classes have no name, and their type must be either a subclass of the named type or an implementer of the named interface.
- ❑ An anonymous inner class is always created as part of a statement, so don't forget to close the statement after the class definition, with a curly brace. This is one of the rare times you'll see a curly brace followed by a semicolon in Java.
- ❑ Because of polymorphism, the only methods you can call on an anonymous inner class reference are those defined in the reference variable class (or interface), even though the anonymous class is really a subclass or implementer of the reference variable type.
- ❑ An anonymous inner class can extend one subclass, *or* implement one interface. Unlike non-anonymous classes (inner or otherwise), an anonymous inner class cannot do both. In other words, it cannot both extend a class *and* implement an interface, nor can it implement more than one interface.
- ❑ An argument-local inner class is declared, defined, and automatically instantiated as part of a method invocation. The key to remember is that the class is being defined within a method argument, so the syntax will end the class definition with a curly brace, followed by a closing parenthesis to end the method call, followed by a semicolon to end the statement:

```
} ) ;
```

Static Nested Classes

- ❑ Static nested classes are inner classes marked with the `static` modifier.
- ❑ Technically, a static nested class is *not* an inner class, but instead is considered a top-level nested class.
- ❑ Because the nested class is static, it does not share any special relationship with an instance of the outer class. In fact, you don't need an instance of the outer class to instantiate a static nested class.
- ❑ Instantiating a static nested class requires using both the outer and nested class names as follows:

```
BigOuter.Nested n = new BigOuter.Nested();
```
- ❑ A static nested class cannot access nonstatic members of the outer class, since it does not have an implicit reference to any outer instance (in other words, the nested class instance does not get an *outer this* reference).

SELF TEST

The following questions will help you measure your understanding of the dynamic and life-altering material presented in this chapter. Because this chapter spans so many different objectives, the questions here are not organized in specific objective categories. Read all of the choices carefully. Choose all correct answers for each question. Take your time. Breathe.

1. Given the following,

```
public class MyOuter {
    public static class MyInner {
        public static void foo() { }
    }
}
```

which statement, if placed in a class *other* than `MyOuter` or `MyInner`, instantiates an instance of the nested class?

- A. `MyOuter.MyInner m = new MyOuter.MyInner();`
 - B. `MyOuter.MyInner mi = new MyInner();`
 - C. `MyOuter m = new MyOuter();`
`MyOuter.MyInner mi = m.new MyOuter.MyInner();`
 - D. `MyInner mi = new MyOuter.MyInner();`
2. Which two are true about a static nested class?
- A. You must have a reference to an instance of the enclosing class in order to instantiate it.
 - B. It does not have access to nonstatic members of the enclosing class.
 - C. Its variables and methods must be static.
 - D. It can be instantiated using `new MyOuter.MyInner();`.
 - E. It must extend the enclosing class.
3. Which constructs an anonymous inner class instance?
- A. `Runnable r = new Runnable() { };`
 - B. `Runnable r = new Runnable(public void run() { });`
 - C. `Runnable r = new Runnable { public void run() {}};`
 - D. `Runnable r = new Runnable() {public void run{}};`
 - E. `System.out.println(new Runnable() {public void run() {}});`
 - F. `System.out.println(new Runnable(public void run() { }));`

4. Given the following,

```
class Boo {
    Boo(String s) { }
    Boo() { }
}
class Bar extends Boo {
    Bar() { }
    Bar(String s) {super(s);}
    void zoo() {
        // insert code here
    }
}
```

which two create an anonymous inner class from within class Bar? (Choose two.)

- A. `Boo f = new Boo(24) { };`
- B. `Boo f = new Bar() { };`
- C. `Boo f = new Boo() {String s;};`
- D. `Bar f = new Boo(String s) { };`
- E. `Boo f = new Boo.Bar(String s) { };`

5. Given the following,

```
1.class Foo {
2.    class Bar{ }
3.}
4.class Test {
5.    public static void main (String [] args) {
6.        Foo f = new Foo();
7.        // Insert code here
8.    }
9.}
```

which statement, inserted at line 5, creates an instance of Bar?

- A. `Foo.Bar b = new Foo.Bar();`
- B. `Foo.Bar b = f.new Bar();`
- C. `Bar b = new f.Bar();`
- D. `Bar b = f.new Bar();`
- E. `Foo.Bar b = new f.Bar();`

6. Which two are true about a method-local inner class?
- A. It must be marked `final`.
 - B. It can be marked `abstract`.
 - C. It can be marked `public`.
 - D. It can be marked `static`.
 - E. It can access private members of the enclosing class.
7. Which is true about an anonymous inner class?
- A. It can extend exactly one class and implement exactly one interface.
 - B. It can extend exactly one class and can implement multiple interfaces.
 - C. It can extend exactly one class or implement exactly one interface.
 - D. It can implement multiple interfaces regardless of whether it also extends a class.
 - E. It can implement multiple interfaces if it does not extend a class.
8. Given the following,

```
public class Foo {
    Foo() {System.out.print("foo");}
    class Bar{
        Bar() {System.out.print("bar");}
        public void go() {System.out.print("hi");}
    }
    public static void main (String [] args) {
        Foo f = new Foo();
        f.makeBar();
    }
    void makeBar() {
        (new Bar() {}).go();
    }
}
```

what is the result?

- A. Compilation fails.
- B. An error occurs at runtime.
- C. foobarhi
- D. barhi
- E. hi

9. Given the following,

```

1. public class TestObj {
2.     public static void main (String [] args) {
3.         Object o = new Object() {
4.             public boolean equals(Object obj) {
5.                 return true;
6.             }
7.         }
8.         System.out.println(o.equals("Fred"));
9.     }
10. }
```

what is the result?

- A. An exception occurs at runtime.
- B. true
- C. fred
- D. Compilation fails because of an error on line 3.
- E. Compilation fails because of an error on line 4.
- F. Compilation fails because of an error on line 8.
- G. Compilation fails because of an error on a line other than 3, 4, or 8.

10. Given the following,

```

1. public class HorseTest {
2.     public static void main (String [] args) {
3.         class Horse {
4.             public String name;
5.             public Horse(String s) {
6.                 name = s;
7.             }
8.         }
9.         Object obj = new Horse("Zippo");
10.        Horse h = (Horse) obj;
11.        System.out.println(h.name);
12.    }
13. }
```

what is the result?

- A. An exception occurs at runtime at line 10.
- B. Zippo
- C. Compilation fails because of an error on line 3.

- D. Compilation fails because of an error on line 9.
- E. Compilation fails because of an error on line 10.
- F. Compilation fails because of an error on line 11.

11. Given the following,

```
1. public class HorseTest {
2.     public static void main (String [] args) {
3.         class Horse {
4.             public String name;
5.             public Horse(String s) {
6.                 name = s;
7.             }
8.         }
9.         Object obj = new Horse("Zippo");
10.        System.out.println(obj.name);
11.    }
12. }
```

what is the result?

- A. An exception occurs at runtime at line 10.
- B. Zippo
- C. Compilation fails because of an error on line 3.
- D. Compilation fails because of an error on line 9.
- E. Compilation fails because of an error on line 10.

12. Given the following,

```
public abstract class AbstractTest {
    public int getNum() {
        return 45;
    }
    public abstract class Bar {
        public int getNum() {
            return 38;
        }
    }
    public static void main (String [] args) {
        AbstractTest t = new AbstractTest() {
            public int getNum() {
                return 22;
            }
        };
    }
};
```

```
        AbstractTest.Bar f = t.new Bar() {
            public int getNum() {
                return 57;
            }
        };
        System.out.println(f.getNum() + " " + t.getNum());
    }
}
```

what is the result?

- A. 57 22
- B. 45 38
- C. 45 57
- D. An exception occurs at runtime.
- E. Compilation fails.

SELF TEST ANSWERS

- A.** MyInner is a static nested class, so it must be instantiated using the fully-scoped name of MyOuter.MyInner.

B is incorrect because it doesn't use the enclosing name in the new. **C** is incorrect because it uses incorrect syntax. When you instantiate a nested class by invoking new on an instance of the enclosing class, you do not use the enclosing name. The difference between **A** and **C** is that **C** is calling new on an instance of the enclosing class rather than just new by itself. **D** is incorrect because it doesn't use the enclosing class name in the variable declaration.
- B** and **D.** **B** is correct because a static nested class is not tied to an instance of the enclosing class, and thus can't access the nonstatic members of the class (just as a static method can't access nonstatic members of a class). **D** uses the correct syntax for instantiating a static nested class.

A is incorrect because static nested classes do not need (and can't use) a reference to an instance of the enclosing class. **C** is incorrect because static nested classes can declare and define nonstatic members. **E** is wrong because...it just is. There's no rule that says an inner or nested class has to extend anything.
- E** is correct. It defines an anonymous inner class instance, which also means it creates an instance of that new anonymous class at the same time. The anonymous class is an implementer of the Runnable interface, so it must override the run () method of Runnable.

A is incorrect because it doesn't override the run () method, so it violates the rules of interface implementation. **B**, **C**, and **D** use incorrect syntax.
- B** and **C.** **B** is correct because anonymous inner classes are no different from any other class when it comes to polymorphism. That means you are always allowed to declare a reference variable of the superclass type and have that reference variable refer to an instance of a subclass type, which in this case is an anonymous subclass of Bar. Since Bar is a subclass of Boo, it all works. **C** uses correct syntax for creating an instance of Boo.

A is incorrect because it passes an *int* to the Boo constructor, and there is no matching constructor in the Boo class. **D** is incorrect because it violates the rules of polymorphism—you cannot refer to a superclass type using a reference variable declared as the subclass type. The superclass is not guaranteed to have everything the subclass has. **E** uses incorrect syntax.

5. **B** is correct because the syntax is correct—using both names (the enclosing class and the inner class) in the reference declaration, then using a reference to the enclosing class to invoke `new` on the inner class.
- A, C, D, and E** all use incorrect syntax. **A** is incorrect because it doesn't use a reference to the enclosing class, and also because it includes both names in the `new`. **C** is incorrect because it doesn't use the enclosing class name in the reference variable declaration, and because the new syntax is wrong. **D** is incorrect because it doesn't use the enclosing class name in the reference variable declaration. **E** is incorrect because the `new` syntax is wrong.
6. **B and E**. **B** is correct because a method-local inner class can be abstract, although it means a subclass of the inner class must be created if the abstract class is to be used (so an abstract method-local inner class is probably not useful). **E** is correct because a method-local inner class works like any other inner class—it has a special relationship to an instance of the enclosing class, thus it can access all members of the enclosing class.
- A** is incorrect because a method-local inner class does not have to be declared `final` (although it is legal to do so). **C and D** are incorrect because a method-local inner class cannot be made `public` (remember—you cannot mark any local variables as `public`), or `static`.
7. **C** is correct because the syntax of an anonymous inner class allows for only one named type after the `new`, and that type must be either a single interface (in which case the anonymous class implements that one interface) or a single class (in which case the anonymous class extends that one class).
- A, B, D, and E** are all incorrect because they don't follow the syntax rules described in the response for answer **C**.
8. **C** is correct because first the `Foo` instance is created, which means the `Foo` constructor runs and prints "foo". Next, the `makeBar()` method is invoked which creates a `Bar`, which means the `Bar` constructor runs and prints "bar", and finally the `go()` method is invoked on the new `Bar` instance, which means the `go()` method prints "hi".
- A, C, D, E, and F** are incorrect based on the program logic described above.
9. **G**. This code would be legal if line 7 ended with a semicolon. Remember that line 3 is a statement that doesn't end until line 7, and a statement needs a closing semicolon!
- A, B, C, D, E, and F** are incorrect based on the program logic described above. If the semicolon were added at line 7, then answer **B** would be correct—the program would print "true", the return from the `equals()` method overridden by the anonymous subclass of `Object`.

10. B. The code in the `HorseTest` class is perfectly legal. Line 9 creates an instance of the method-local inner class `Horse`, using a reference variable declared as type `Object`. Line 10 casts the `Horse` object to a `Horse` reference variable, which allows line 11 to compile. If line 10 were removed, the `HorseTest` code would not compile, because class `Object` does not have a *name* variable.
- A, C, D, E, and F are incorrect based on the program logic described above.
11. E. This code is identical to the code in question 10, except the casting statement has been removed. If you use a reference variable of type `Object`, you can access only those members defined in class `Object`.
- A, B, C, and D are incorrect based on the program logic described above and in the previous question.
12. A. You can define an inner class as `abstract`, which means you can instantiate only concrete subclasses of the abstract inner class. The object referenced by the variable *t* is an instance of an anonymous subclass of `AbstractTest`, and the anonymous class overrides the `getNum()` method to return 22. The variable referenced by *f* is an instance of an anonymous subclass of `Bar`, and the anonymous `Bar` subclass also overrides the `getNum()` method (to return 57). Remember that to instantiate a `Bar` instance, we need an instance of the enclosing `AbstractTest` class to tie to the new `Bar` inner class instance. `AbstractTest` can't be instantiated because it's abstract, so we created an anonymous subclass (non-abstract) and then used the instance of that anonymous subclass to tie to the new `Bar` subclass instance.
- B, C, D, E, and F are incorrect based on the program logic described above.

