

Java SOLID Principles

Table of contents

Why SOLID?	1
Single Responsibility Principle (SRP)	2
Open-Closed Principle (OCP)	2
Liskov Substitution Principle (LSP)	4
Interface Segregation Principle (ISP)	6
Dependency Inversion Principle (DIP)	7
Summary	10

Introduced in 2026-04-02-refactoring¹ (Lecture, Week 6, Part 3 — Single Responsibility and Open-Closed) and continued in 2026-04-16-solid-principles-ii² (Lecture, Week 7 — Liskov Substitution, Interface Segregation, Dependency Inversion).

Why SOLID?

Large software tends to become:

- **Rigid** — difficult to change; even small changes require modifications across many parts of the codebase.
- **Fragile** — likely to break when changed, sometimes in parts of the codebase that appear unrelated to the change.

The SOLID principles (S-ingle responsibility, O-pen-closed, L-iskov substitution, I-nterface segregation, D-eendency inversion) are a set of guidelines — not rules — to help prevent rigidity and fragility. Many of the individual ideas predate the acronym; Robert Martin combined them in *Design Principles and Design Patterns* (2000).

¹courses/csse2002/2026-04-02-refactoring.pdf

²courses/csse2002/2026-04-16-solid-principles-ii.pdf

Single Responsibility Principle (SRP)

There should never be more than one reason for a class to change.

Closely related to java-cohesion-and-coupling³'s cohesion: cohesion is about how *similar* a class's functionality is, while SRP is about the *reasons for change*. A “responsibility” is a reason for change — if you can think of more than one motive to change a class, it has more than one responsibility.

```
class Student {
    private String name;
    private List<Course> courses;
    private double gpa;

    public void enrollInCourse(Course course) { courses.add(course); }
    public void calculateGPA() { /* ... */ }
    public void printTranscript() { /* ... */ }
}
```

This `Student` class has (at least) three reasons to change: how enrolment works, how GPA is calculated, and how transcripts are generated. Splitting each concern into its own class gives each one a single responsibility:

```
class Student {
    private String name;
    private List<Course> courses;
    public void enrollInCourse(Course course) { courses.add(course); }
    public List<Course> getCourses() { return courses; }
}
class GPACalculator {
    public double calculateGPA(List<Course> courses) { /* ... */ }
}
class TranscriptPrinter {
    public void printTranscript(Student student) { /* ... */ }
}
```

Open-Closed Principle (OCP)

Components should be **open for extension** but **closed for modification**.

Introduced by Bertrand Meyer in *Object Oriented Software Construction* (1988): “if the open-closed principle is applied well, then further changes are achieved by adding new code, not by changing old code that already works”.

³[courses/csse2002/java-cohesion-and-coupling.pdf](#)

- **Open for extension** — a class's behaviour can be extended to support changing requirements.
- **Closed for modification** — no one may change the behaviour of an existing class; it's closed once it's available for other modules to depend on.

How do we keep a class closed for modification? Information hiding (see java-encapsulation⁴).

Example: pluggable grading systems

A university needs to support multiple grading systems (letter grades, numeric grades). A first attempt branches on a `gradeType` string inside `GPACalculator` — every new grading system means *editing* this method (violates OCP):

```
class GPACalculator {
    public double calculateGPA(List<Course> courses, String gradeType) {
        double totalPoints = 0;
        if (gradeType.equals("Letter")) {
            for (Course course : courses) {
                totalPoints += convertLetterGradeToPoints(course.getLetterGrade());
            }
        } else if (gradeType.equals("Numeric")) {
            for (Course course : courses) {
                totalPoints += course.getNumericGrade();
            }
        }
        return totalPoints / courses.size();
    }
}
```

Extracting the grading logic behind an interface makes `GPACalculator` open for extension (new `GradingSystem` implementations) without ever touching its own code again:

```
interface GradingSystem {
    double getGradePoints(Course course);
}

class LetterGradingSystem implements GradingSystem {
    @Override
    public double getGradePoints(Course course) { /* convert letter grade to points
        ↪ */ return 0; }
}

class NumericGradingSystem implements GradingSystem {
    @Override
```

⁴courses/csse2002/java-encapsulation.pdf

```

    public double getGradePoints(Course course) { return course.getNumericGrade(); }
}

class GPACalculator {
    private GradingSystem gradingSystem;
    public GPACalculator(GradingSystem gradingSystem) { this.gradingSystem =
        ↪ gradingSystem; }

    public double calculateGPA(List<Course> courses) {
        double totalPoints = 0;
        for (Course course : courses) {
            totalPoints += gradingSystem.getGradePoints(course);
        }
        return totalPoints / courses.size();
    }
}

```

GPACalculator now depends on the `GradingSystem` *abstraction* rather than calculating grade points itself — new grading systems (e.g. pass/fail) can be added just by writing a new class that implements `GradingSystem`, with zero changes to `GPACalculator`'s own code.

Liskov Substitution Principle (LSP)

Subclasses should be substitutable for their parent classes.

Defined by Barbara Liskov in a 1987 keynote: an instance of a subclass type can be used wherever an instance of a superclass type is expected, *without breaking the correctness of the program*:

```

List<Animal> animals = new ArrayList<>();
animals.add(new Cat());
animals.add(new Mouse());
for (Animal animal : animals) {
    animal.eat(new Food());
}

```

Example violation

```

class Student {
    protected List<Course> courses;
    public double calculateGPA() {
        double totalPoints = 0;

```

```

        for (Course course : courses) {
            totalPoints += course.getGradePoints();
        }
        return totalPoints / courses.size();
    }
}
class PostgraduateStudent extends Student {
    @Override
    public double calculateGPA() {
        throw new UnsupportedOperationException("Postgraduates have a different GPA
        ↪ system.");
    }
}

```

Code written against `Student` (e.g. calling `calculateGPA()` on every student in a list) will blow up the moment a `PostgraduateStudent` is substituted in — `PostgraduateStudent` is not actually usable wherever a `Student` is expected, so it violates LSP.

Substitution with contracts

Even when a subclass overrides a method with a different implementation, it must still honour the parent’s contract (see java-specification⁵):

```

class Parent {
    /**
     * @requires x > 0
     * @ensures \result > 'A' && \result < 'Z'
     */
    char f(int x);
}

```

Preconditions — suppose an overriding `Child.f()` can also handle negative numbers. Changing its precondition from `x > 0` to `x != 0` is fine: every input `Parent.f()` accepted (`x > 0`) is still accepted by `Child.f()` (`x != 0`), since $x > 0 \Rightarrow x \neq 0$ — so the precondition became *weaker* (less strict), not stronger. **Preconditions in a subclass must be no stronger than the superclass’s — they may be weaker or stay the same.**

Postconditions — suppose `Child.f()` only ever returns 'K', 'L', or 'M'. A postcondition of `\result > 'J' && \result < 'N'` is fine: every result still satisfies the parent’s promise (`\result > 'J' && \result < 'N' \Rightarrow \result > 'A' && \result < 'Z'`), so the postcondition became *stronger* (more restrictive), which is allowed. **Postconditions in a subclass must be no weaker than the superclass’s — they may be stronger or stay the same.**

⁵[courses/csse2002/java-specification.pdf](https://courses.csse2002/java-specification.pdf)

A subclass must not strengthen preconditions or weaken postconditions — it should accept everything the parent accepts, and still guarantee everything the parent guarantees.

Worked example

```
class One {
    /**
     * @require row != null && col != null && 0 < val && val < 100
     * @ensure \result is the list of all of the positive
     *         integers n < val that appear in either row or col
     */
    public List<Integer> calc(int[] row, int[] col, int val) { ... }
}
class Two extends One {
    /**
     * @require row != null && col != null && 0 < val &&
     *         val < 100 && row only contains positive integers
     * @ensure \result is the list of all of the positive
     *         integers n < val that appear in either row or col
     */
    @Override
    public List<Integer> calc(int[] row, int[] col, int val) { ... }
}
```

Two's precondition adds an extra requirement (*row only contains positive integers*) that One never required — so One's precondition does *not* imply Two's: there are inputs One would accept (a *row* containing a negative number) that Two rejects. This **strengthens** the precondition and violates LSP: code written against One that happens to pass a negative-containing *row* would break if a Two were substituted in.

Homework variant: if Two's postcondition is instead changed to return integers appearing in **both** *row* and *col* (rather than *either*), does this satisfy LSP? Reasoning it through: the guarantee changes from a union to an intersection, which for most inputs returns strictly *fewer* elements than One promised — that's a **weaker** postcondition (some outputs a caller was guaranteed under One are no longer guaranteed under Two), which also violates LSP.

Interface Segregation Principle (ISP)

Many client-specific interfaces are better than one general-purpose interface.

Shrink interfaces to minimize *incidental dependencies* — large interfaces should be split into smaller ones, so implementing/using classes only need to be concerned about the methods that actually interest them.

```
class Z {
    public void a() {...}
    public void b() {...}
    public void c() {...}
    public void d() {...}
    public void e() {...}
}
// A only uses a() and b(), B only uses c(), C only uses d() and e() ...
// ...but all three classes still depend on the whole of Z.
```

This mirrors java-cohesion-and-coupling⁶'s stamp coupling at the interface level: despite each client using only a small subset of Z, every client depends on *all* of Z — so a change to any part of Z risks affecting (and forcing a recompile/redeploy of) every client, even ones that never used the changed part. Splitting Z's surface into focused interfaces removes those incidental dependencies:

```
interface AI { void a(); void b(); }
interface BI { void c(); }
interface CI { void d(); void e(); }
// Z implements AI, BI, CI; A depends only on AI, B only on BI, C only on CI.
```

Dependency Inversion Principle (DIP)

Depend upon abstractions. Do not depend upon concretions.

Entities must depend on abstractions, not concretions: high-level modules must not depend on low-level modules — both should depend on abstractions. If A depends directly on the concrete class B, changes to B's implementation can propagate to A; if A instead depends only on an abstraction (BSpec) that B implements, changes inside B are far less likely to affect A, as long as BSpec itself stays the same.

Even a small step towards this helps — programming against the `List` interface instead of the concrete `ArrayList` type means `Student` no longer depends on which list implementation the caller chose:

⁶[courses/csse2002/java-cohesion-and-coupling.pdf](https://courses.csse2002/java-cohesion-and-coupling.pdf)

```
class Student {
    private List<Course> courses; // not ArrayList<Course>
    public Student(List<Course> courses) { this.courses = courses; }
    public double calculateTotalGradePoints() { ... }
}
```

What makes a good dependency? (stability and exposure)

A class A has a **dependency** on class B if A refers to B in code (an `import B;`, or, in the same package, simply referring to B anywhere). Not all dependencies are equally risky — two properties determine how good/bad one is:

- **Stability** — how likely the dependency is to change. `ArrayList` is a good dependency because it's very unlikely to change; treat classes *you* write as unstable until proven otherwise, and generally assume interfaces are more stable than concrete classes (unless they're poorly designed and change often).
- **Exposure** — how much of the depending class is entangled with it. A dependency only used in a constructor is better than one used across every method, since less of the class would need to change if that dependency changed.

Two forms of DIP

Minimising concrete dependencies (the simple form) — change a field's compile-time type from a concrete class to whatever interface it already implements, wherever one already exists:

```
- HardwoodFloor placedOn;
+ Floor placedOn;
```

This only helps where an abstraction already exists, though. The **proper form of DIP** goes further:

Remove dependencies on low-level components from high-level components — make *both* depend on an abstraction.

“High-level” vs “low-level” isn't a strict distinction: high-level components implement application logic (e.g. a browser's back/forward navigation history); low-level components do the grunt work that logic depends on (e.g. actually requesting a webpage's data). The goal is for high-level components to talk to low-level ones *exclusively* through an abstraction, so a low-level component changing doesn't force the high-level component depending on it to change too. Applied to a `Desk` class depending on a concrete `LogitechKeyboard` with no existing interface:

1. **Create an abstraction** of the low-level component: `interface Keyboard { void type(char key); ... }`.
2. **Modify the low-level component** to implement it: `class LogitechKeyboard implements Keyboard`.
3. **Minimise concrete dependencies** in the high-level component: `Keyboard keyboard = new LogitechKeyboard();`.

Dependency Injection

Dependency Injection (DI) is a design pattern that implements DIP by injecting a class's dependencies into it (e.g. via the constructor) instead of having the class create them internally, promoting loose coupling (see [java-cohesion-and-coupling⁷](#)):

```
// Before: AlertService is tightly coupled to EmailSender, and creates its own
↳ dependency.
class AlertService {
    public void sendAlert(String msg) {
        EmailSender sender = new EmailSender();
        sender.send(msg);
    }
}
```

```
// After: AlertService depends on the MessageSender abstraction, injected via the
↳ constructor.
interface MessageSender {
    void send(String msg);
}
class EmailSender implements MessageSender {
    @Override
    public void send(String msg) { System.out.println("Sending EMAIL: " + msg); }
}
class SmsSender implements MessageSender {
    @Override
    public void send(String msg) { System.out.println("Sending SMS: " + msg); }
}
class AlertService {
    private MessageSender sender;
    public AlertService(MessageSender sender) { this.sender = sender; }
    public void sendAlert(String msg) { sender.send(msg); }
}
```

`AlertService` can now send alerts via email, SMS, or any future `MessageSender` implementation, without ever changing its own code.

⁷[courses/csse2002/java-cohesion-and-coupling.pdf](#)

A constructor isn't the only injection point — a setter is appropriate when the dependency is likely to change over the object's lifetime, e.g. `public changeSender(MessageSender sender) { this.sender = sender; }`. Eventually *something* (typically an entry-point method like `main`) has to construct the concrete dependencies before injecting them; DI frameworks let you defer that construction to runtime via the framework's own configuration instead of in code, but unless a project needs multiple *levels* of injected dependencies, good design alone can avoid the extra conceptual overhead of adopting one.

Summary

- **Single Responsibility** — one reason to change per class.
- **Open-Closed** — open for extension, closed for modification.
- **Liskov Substitution** — subclasses substitutable for their parents; no stronger preconditions, no weaker postconditions.
- **Interface Segregation** — many small, client-specific interfaces beat one large general-purpose interface.
- **Dependency Inversion** — depend on abstractions, not concretions; inject dependencies rather than constructing them internally.