47 · Scalable Applications

Flutter Practice Notes

Contents

47 · Scalable Applications

Introduction

This module covers building Flutter apps that **scale** across four dimensions — **codebase** (hundreds of files/features), **team** (many contributors), **runtime** (startup/memory/frame budget under growth), and **features** (adding capability without regressions) — and the practices that keep all four healthy: code organization + governance at scale, team topologies + ownership, performance-at-scale, design systems + shared packages, and **technical-debt/evolution** management. It synthesizes the architecture band (Clean/MVVM/feature-first/modular/DDD — Modules 40–46) into a coherent "how to grow" playbook, capped by a capstone.

Why this module exists

Patterns that work at 5 screens break at 500: a codebase becomes unnavigable, a team steps on itself, startup/memory degrade, and every new feature risks regressions elsewhere. Scaling isn't one technique — it's **managing the tradeoffs across all four dimensions simultaneously**, with the right structure, governance, performance discipline, and debt strategy for the app's current stage. This module teaches the dimensions, the levers, and — crucially — **right-sizing** them to where the app actually is.

Module map

#	File	Topic	Level
1	01_scaling_dimensions.md	The four dimensions of scale; right-sizing	🔴
2	02_codebase_and_team_scaling.md	Organization, conventions, governance, ownership, design systems	🔴
3	03_performance_and_runtime_scaling.md	Startup, memory, frame budget, deferred loading at scale	🔴
4	04_technical_debt_and_evolution.md	Managing debt, refactoring, deprecation, migration at scale	🟡
5	05_scalable_integration.md	Capstone: a scaling playbook	🔴

> **Cross-references:** Architecture backbone: 40/43/44/45/46. Performance: Module 21. Design systems: Module 25. CI/CD: Module 50. Monitoring: Module 52. Testing: Module 49. System design: Module 48.

Prerequisites

The architecture band (40–46), 21 Performance, 49 Testing, 50 CI CD.

What you'll be able to do after this module

- Identify the four scaling dimensions and diagnose which is under pressure.
- Organize a large codebase with conventions, governance, ownership, and a design system.
- Keep startup, memory, and frame budget healthy as features multiply.
- Manage technical debt, refactoring, deprecation, and large migrations deliberately.
- Assemble a right-sized scaling playbook for an app's current stage.

Capstone

Scaling playbook: For a hypothetical growing app, produce a playbook covering: modular/feature-first structure + governance (conventions, lints, CODEOWNERS), a shared design-system package, a performance budget (startup/memory/frame) with deferred loading, a tech-debt register + refactoring/deprecation process, and CI/monitoring hooks — each right-sized to the app's stage with explicit "do this now / defer this" calls.

Summary

Scaling means keeping codebase, team, runtime, and feature growth healthy together — via modular/feature-first structure + governance, team ownership + a design system, performance budgets + deferred loading, and deliberate tech-debt/evolution management. There's no single fix; you right-size the levers to the app's stage and manage the tradeoffs continuously.

The Four Dimensions of Scale

> "Scaling" isn't one thing — a Flutter app scales along **four independent dimensions**, each with its own pressure points and levers: **codebase** (files/features → navigability, build time, coupling), **team** (contributors → coordination, ownership, merge conflicts), **runtime** (features/data/users → startup, memory, frame budget), and **features** (capability growth → regression risk, consistency). The scaling mistake is treating them as one: you over-engineer one dimension while another silently degrades. The discipline is to **diagnose which dimension is under pressure** and apply the **right-sized lever** — not maximal architecture everywhere.

Introduction

This file frames scaling as four distinct dimensions, their symptoms, their levers, and — most importantly — **right-sizing**: applying just enough of each lever for the app's current stage. It's the diagnostic lens for the rest of the module.

Why this concept exists

Teams conflate scaling with "more architecture," then either gold-plate a small app (wasted effort) or bolt on structure reactively when one dimension is already on fire. Naming the four dimensions lets you diagnose the *actual* pressure and respond proportionally — the difference between deliberate growth and thrash.

Real-world analogy

Scaling a city has independent dimensions: **road network** (codebase — navigability/throughput), **workforce/governance** (team — coordination), **utilities capacity** (runtime — power/water under load), and **new districts** (features — growth without breaking existing ones). A city that only widens roads while the power grid fails hasn't scaled. You **monitor each system, find the bottleneck, and invest there** — not pour concrete everywhere.

Internal Working

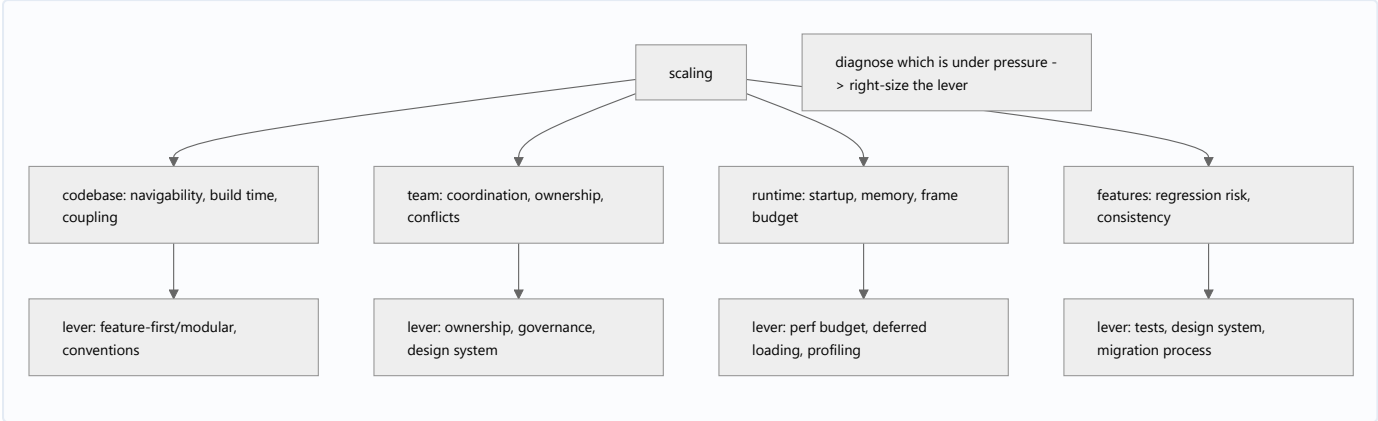

- **Codebase scale** (more files/features/lines):
  - *Symptoms*: hard to find/understand code, slow builds, tangled dependencies, wide-blast-radius changes.
  - *Levers*: **feature-first** → **modular** structure (44/45), conventions, import boundaries/lints, incremental builds.
- **Team scale** (more contributors):
  - *Symptoms*: merge conflicts, unclear ownership, inconsistent patterns, coordination overhead, onboarding pain.
  - *Levers*: **ownership** (CODEOWNERS/packages), **governance** (conventions, lints, review), a **design system** (consistency), parallelizable module boundaries, docs/onboarding.
- **Runtime scale** (more features/data/users on-device):
  - *Symptoms*: slow startup, high memory, jank on big lists/many features, large app size.
  - *Levers*: **performance budget** (startup/memory/frame), deferred/lazy loading, profiling, list virtualization, scoped rebuilds (Module 21).
- **Feature scale** (adding capability safely):
  - *Symptoms*: new features break old ones, inconsistent UX, fear of changing shared code.
  - *Levers*: **tests** (regression safety — Module 49), **design system** (consistent UX), **stable contracts** + **migration process**, feature flags.
- **Dimensions are independent**: an app can be codebase-heavy but low-team (solo dev with 300 files), or high-team but simple-runtime. **Optimize the dimension under pressure**, not all at once.
- **Right-sizing (the core skill)**: apply the **minimum lever** that relieves the current bottleneck; **defer** the rest. A 3-screen prototype needs none of this; a 50-feature multi-team app needs most. **Grow into** structure (Module 44) rather than pre-building it.
- **Diagnosis first**: measure/observe the symptom (build times, conflict rate, frame stats, regression frequency) to identify the pressured dimension, then invest there — data over

dogma.

- **The dimensions interact**: modularization (codebase) also helps team (ownership) + build (runtime tooling); a design system helps team + feature consistency — pick levers that relieve multiple dimensions when possible.

Memory Representation

Not runtime — a **diagnostic model**: four dimensions × current-pressure × available levers. The mental artifact is "which dimension is the bottleneck now, and what's the smallest lever?"

Compiler Behavior

Codebase-scale levers (modular boundaries, lints) are compile/build-enforced; the others are process/runtime concerns.

Runtime Behavior

Only the runtime dimension has direct runtime effects (startup/memory/frame); the others affect build/dev/team velocity.

Flutter Engine Behavior

Runtime-scale levers interact with the engine (rebuilds, rasterization, memory — Module 21); the rest don't.

Dart VM Behavior

Codebase modularization enables incremental builds; runtime scaling touches GC/isolates; others are org/process.

Examples

```
Diagnose-then-lever (right-sizing) examples:
  Symptom: builds take 8 min, changes ripple widely  -> dimension: CODEBASE -> lever: modularize (Module 45)
  Symptom: constant merge conflicts, unclear owners   -> dimension: TEAM     -> lever: CODEOWNERS + module boundaries + design system
  Symptom: 3s cold start, jank on the feed            -> dimension: RUNTIME -> lever: perf budget + deferred loading + list virtualization (Module 21)
  Symptom: every release breaks something old         -> dimension: FEATURE -> lever: tests + stable contracts + migration process (Module 49)

Anti-pattern: a 4-screen prototype adopting melos monorepo + DDD -> over-engineering (no dimension under pressure)
```

Diagrams

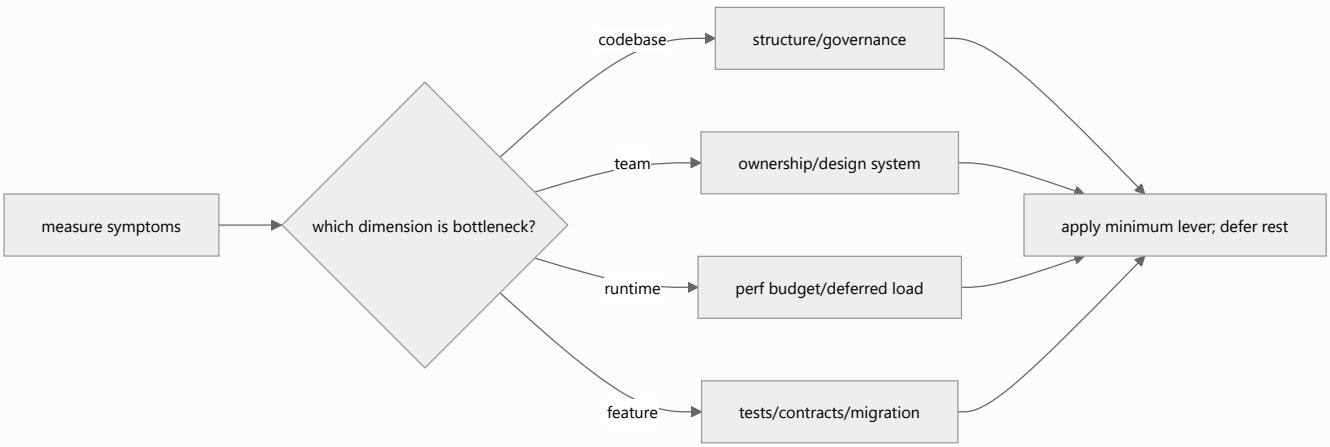

Common Mistakes

Mistake	Why it's wrong	Fix
Treating scale as one thing	Over-build one, neglect another	Diagnose per dimension
Pre-building max architecture on a small app	Wasted effort, slower delivery	Right-size; grow into structure
Reactive scaling (fix when on fire)	Painful, risky	Monitor symptoms; invest proactively-but-proportionally
Dogma over data	Wrong lever for the bottleneck	Measure symptoms, then act
Ignoring dimension interactions	Miss multi-benefit levers	Prefer levers that relieve several
Scaling runtime by adding architecture	Wrong lever	Runtime → profiling/budget, not folders
Never scaling ("it works")	Slow collapse at growth	Watch for symptoms; invest in time

Best Practices

- Treat scale as **four independent dimensions** (codebase, team, runtime, feature); **diagnose which is under pressure** from real symptoms before acting.
- Apply the **minimum lever** that relieves the current bottleneck; **defer** the rest; **grow into** structure rather than pre-building it.
- Prefer levers that **relieve multiple dimensions** (modularization → codebase + team + build; design system → team + feature).
- **Measure** (build times, conflict rate, frame/startup/memory stats, regression frequency) to guide investment — **data over dogma**; re-diagnose as the app grows.

Performance

Only the runtime dimension is about runtime performance; the others are about **developer/team performance** (velocity, safety, coordination). Right-sizing is itself a performance discipline — avoiding wasted effort on non-bottleneck dimensions.

Advantages / Disadvantages

- **+** (Right-sizing per dimension) targeted investment, sustained velocity, no over/under-engineering, healthy growth on all fronts.
- **−** Requires ongoing diagnosis/measurement + judgment; easy to misjudge the bottleneck; no one-size formula.

Interview Questions

1. 🟢 **What are the four dimensions of scaling an app?** — Codebase, team, runtime, and feature — each with distinct symptoms and levers.
2. 🟢 **Why is treating scaling as one thing a mistake?** — You over-engineer one dimension while another silently degrades; each needs its own diagnosis + lever.
3. 🟡 **How do you decide what to invest in?** — Diagnose which dimension is under pressure from real symptoms (build times, conflicts, frame/startup stats, regressions), then apply the minimum lever.
4. 🟡 **Give a symptom→dimension→lever example.** — Slow builds + wide-ripple changes → codebase → modularize; constant conflicts + unclear owners → team → CODEOWNERS + boundaries + design system.
5. 🟡 **What does right-sizing mean here?** — Applying just enough of each lever for the app's current stage, deferring the rest — growing into structure, not pre-building it.
6. 🔴 **Which levers relieve multiple dimensions?** — Modularization (codebase + team + build), design system (team + feature consistency) — prefer these.
7. 🔴 **How is scaling the runtime dimension different from the others?** — It's the only one about runtime perf (startup/memory/frame → profiling/budget/deferred loading); the others are dev/team/process levers.

Senior Engineer Tips

- Diagnose before you architect: name the pressured dimension from real symptoms, then reach for the smallest lever — most "we need to re-architect" urges are one dimension, fixable with one targeted change.
- Grow into structure; a 4-screen app adopting melos + DDD is as much a scaling failure as a 50-feature app with everything in one file.
- Favor levers with cross-dimension payoff (modular boundaries, a design system) and keep measuring — the bottleneck moves as the app grows, so re-diagnose periodically.

Architect Perspective

Scaling is portfolio management across four dimensions: you continuously diagnose which is the bottleneck and invest the minimum lever there, using the architecture band's tools (feature-first/modular for codebase+team, performance discipline for runtime, tests+contracts for features) proportionally. The architect's value is judgment — right-sizing over dogma, cross-dimension levers, and data-driven re-diagnosis — so the app grows sustainably instead of over-engineering one axis while another collapses (02_codebase_and_team_scaling.md, 03_performance_and_runtime_scaling.md, Module 21).

Summary

- Apps scale on four independent dimensions: codebase, team, runtime, feature — each with its own symptoms and levers.
- Diagnose the bottleneck from real symptoms, apply the minimum lever, defer the rest; grow into structure; prefer multi-dimension levers.
- Only the runtime dimension is about runtime perf; the rest are dev/team/process — right-size with data, not dogma.

Revision Notes

- Dimensions: codebase (navigability/build/coupling → feature-first/modular/lints), team (conflicts/ownership → CODEOWNERS/governance/design system), runtime (startup/memory/frame → perf budget/deferred load/profiling), feature (regressions/consistency → tests/contracts/migration).
- Independent → diagnose the bottleneck (measure symptoms), apply minimum lever, defer rest, grow into structure; prefer multi-dimension levers (modular, design system).
- Right-sizing = data over dogma; only runtime dimension = runtime perf; re-diagnose as app grows.

Practice Questions

1. Name the four dimensions and a symptom of each.
2. How do you decide which dimension to invest in?
3. Why is pre-building maximal architecture a scaling failure?

Coding Questions

1. Given a set of symptoms, map each to a dimension + a right-sized lever.
2. Identify a lever that relieves multiple dimensions and explain how.
3. Justify deferring a scaling investment for a given app stage.

Mini Project

Scale diagnosis (Flutter): For three app scenarios (a prototype, a growing single-team app, a large multi-team app), diagnose which scaling dimension(s) are under pressure from given symptoms, choose right-sized levers (deferring the rest), and identify any multi-dimension levers. Acceptance: four dimensions applied; symptoms → dimension → minimum-lever mapping per scenario; right-sizing (defer non-bottlenecks) justified; multi-dimension levers noted; data-over-dogma stance.

Codebase & Team Scaling

> As code and contributors grow, two dimensions must be managed together: **codebase** (keep it navigable, low-coupling, buildable) and **team** (keep contributors coordinated, unblocked, consistent). The levers overlap: **feature-first → modular structure** (cohesion + enforced boundaries), **clear ownership** (packages/CODEOWNERS = one team per unit), **governance** (shared conventions, lints/analysis, PR standards, docs), and a **shared design system** (one consistent UI vocabulary). The goal is **parallel work without collisions or drift** — teams own modules, boundaries prevent stepping on each other, and conventions keep the codebase coherent.

Introduction

This file covers scaling the codebase and team dimensions jointly: structure for navigability + boundaries, ownership mapping, governance to prevent drift, and a design system for UI consistency. It applies the feature-first/modular structure (44/45) to the human/organizational problem.

Why this concept exists

A big codebase with many contributors and no structure becomes unnavigable, conflict-ridden, and inconsistent — each dev invents patterns, changes ripple unpredictably, and no one owns anything. Structure + ownership + governance + a design system turn a growing crowd into coordinated teams working in parallel on cohesive, boundaried units with a shared style.

Real-world analogy

It's running a **large construction firm**: work is divided into **crews owning specific buildings** (module ownership), with **blueprints/building codes everyone follows** (conventions/lints), **inspections before work is accepted** (PR review/CI), and a **shared catalog of standard fixtures** (design system) so every building looks coherent. Without these, 40 workers on one site produce chaos, mismatched fittings, and constant collisions.

Internal Working

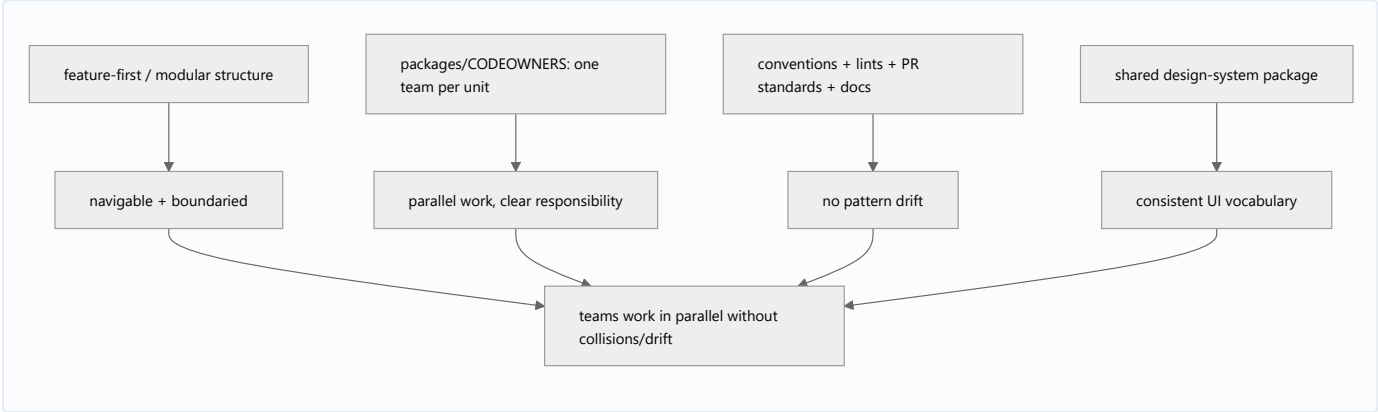

- **Structure (codebase)**: **feature-first** organization (cohesive vertical slices) graduating to **modular packages** with **enforced boundaries** as size/teams grow (44/45). Cohesion makes code findable; boundaries limit blast radius + enable parallel work + incremental builds.
- **Ownership (team)**: map **modules/packages to owning teams** (CODEOWNERS); a unit has **one clear owner** responsible for its API, quality, and reviews. This aligns Conway's Law deliberately (module boundaries ≈ team boundaries) and removes "everyone and no one" responsibility. Cross-team interaction goes through **contracts** (Module 45).
- **Governance (team + codebase)** — prevents drift as contributors multiply:
  - **Conventions**: agreed patterns (architecture, naming, folder layout, state management) documented and applied uniformly.
  - **Automated enforcement**: **lints/analysis** (custom + `analysis_options.yaml`), **import-boundary rules**, formatters, pre-commit hooks, **CI gates** (Module 50) — machines enforce so reviews focus on substance.
  - **PR standards + review**: size limits, required reviewers (via CODEOWNERS), templates, definition-of-done.
  - **Docs/onboarding**: architecture guide, ADRs (architecture decision records), a runnable example slice as the template, onboarding docs — so new devs ramp fast and follow patterns.
- **Design system (team + feature consistency)**: a **shared `design_system` / `core/ui` package** of tokens (colors/spacing/typography) + reusable components (buttons, inputs, cards) + theming, owned by a platform/design team. It gives every feature **one consistent UI vocabulary**, prevents each team reinventing widgets, and makes global restyling one change. It's the UI analog of shared `core` (Module 25/Module 44).
- **The joint goal**: **parallel work without collisions (boundaries + ownership) or drift (governance + design system)** — the codebase stays coherent and buildable as the team grows.
- **Right-sizing**: solo/small → light conventions + feature-first folders; growing/multi-team → modular packages + CODEOWNERS + strict governance + a design system. Introduce as the

team/codebase pressure appears (01_scaling_dimensions.md).

Memory Representation

Not runtime — an **org + repo structure**: modules/packages ↔ owning teams, a conventions/ADR corpus, automated-enforcement config, and a design-system package. The invariant: cohesive owned units with enforced boundaries + shared conventions/UI.

Compiler Behavior

Lints/analysis + import-boundary rules (→ package boundaries) enforce conventions at build time; the design-system package's public API constrains UI usage. Governance becomes partly compiler-enforced, not just review-enforced.

Runtime Behavior

Not applicable — these are dev/build/org concerns. (A shared design system does yield consistent runtime UI.)

Flutter Engine Behavior

Design-system components render like any widgets; consistency reduces bespoke, error-prone UI.

Dart VM Behavior

Modular packages enable incremental builds (codebase-scale benefit — Module 45).

Examples

```
Governance stack (automated first, review second):
  analysis_options.yaml   -> lints + custom rules (enforce patterns)
  import boundary rules    -> features -> core/contracts only (no feature->feature)
  formatter + pre-commit   -> consistent style, no bikeshedding
  CI gates                 -> analyze + test (scoped) must pass (Module 50)
  CODEOWNERS               -> per-package owners; required reviewers
  ADRs + architecture doc  -> decisions recorded; example slice = the template
```

```
# CODEOWNERS — modules mapped to owning teams (Conway's Law, on purpose)
/packages/design_system/   @team-platform
/packages/feature_cart/    @team-commerce
/packages/feature_auth/    @team-identity
/packages/contracts/       @team-architecture   # cross-team contracts gated
```

```
// Design system: one UI vocabulary every feature reuses (owned centrally)
// package: design_system
class AppButton extends StatelessWidget { /* token-driven, themed, consistent */ }
class AppTextField extends StatelessWidget { /* ... */ }
// features import design_system for UI; they don't reinvent buttons/inputs.
```

Diagrams

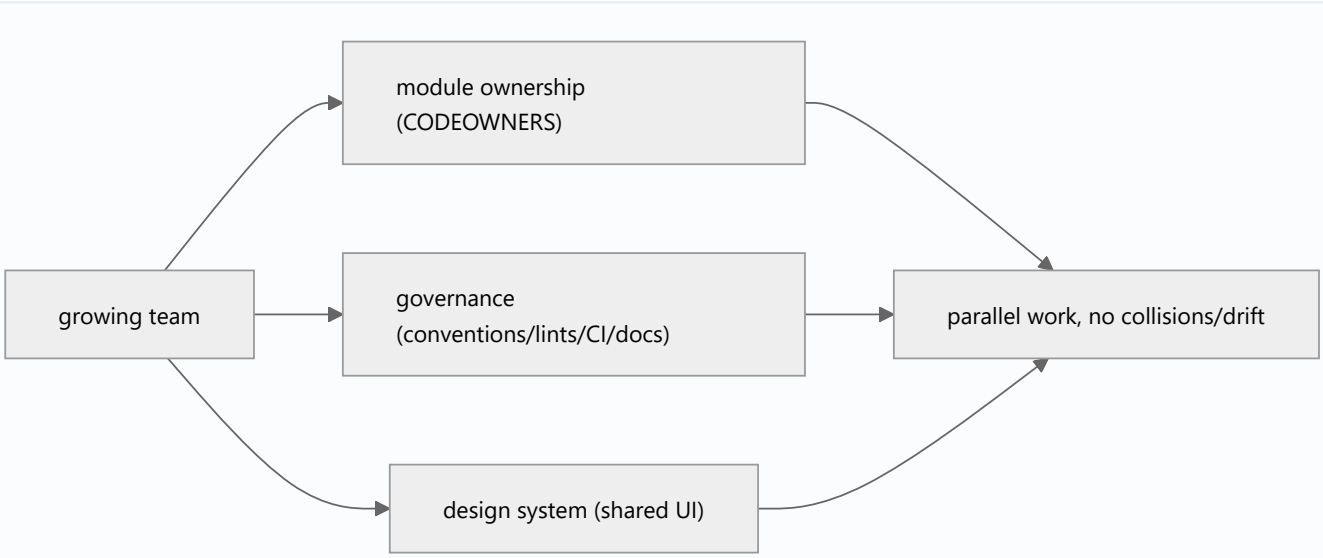

Common Mistakes

Mistake	Why it's a problem	Fix
No ownership (shared everything)	Diffuse responsibility, conflicts	Module/package owners (CODEOWNERS)
Conventions unenforced (docs only)	Drift as team grows	Automate (lints/CI/import rules)
Each team reinventing UI	Inconsistent UX, duplication	Shared design-system package
No boundaries between teams' code	Collisions, ripple changes	Modular boundaries + contracts
Manual review for style/patterns	Slow, inconsistent	Automate enforcement; review substance
No onboarding docs/ADRs/example	Slow ramp, pattern divergence	Architecture guide + ADRs + template slice
Heavy governance on a tiny team	Overhead	Right-size (light for small teams)

Best Practices

- Use **feature-first → modular** structure with **enforced boundaries**; map **modules/packages to owning teams** (CODEOWNERS) so each unit has one clear owner.
- **Automate governance** (lints, import-boundary rules, formatter, CI gates) so patterns can't drift; reserve **review** for substance; record decisions in **ADRs** + an architecture guide + an **example slice** as the template.
- Provide a **shared design-system package** (tokens + components + theming) so features share one UI vocabulary (owned centrally); cross-team interaction via **contracts**.
- **Right-size** governance/ownership to team size; grow structure as codebase/team pressure appears (01_scaling_dimensions.md).

Performance

Developer-velocity performance: cohesion + boundaries speed navigation + incremental builds; automated governance removes review bottlenecks; a design system removes UI reinvention; ownership removes coordination stalls. No runtime cost (design system yields consistent runtime UI).

Advantages / Disadvantages

- **+** Parallel work without collisions/drift, clear ownership, consistent UI, fast onboarding, incremental builds, coherent large codebase.
- **−** Governance/design-system setup + maintenance, ownership boundaries to define + respect, overhead if applied to a tiny team, cross-team contract discipline.

Interview Questions

1. 🟢 **How do you keep a large codebase navigable with many contributors?** — Feature-first → modular structure (cohesion + enforced boundaries), module/package ownership, governance (conventions/lints/CI), and a shared design system.
2. 🟢 **How does ownership scale a team?** — Map modules/packages to owning teams (CODEOWNERS) so each unit has one responsible owner, enabling parallel work + clear accountability; interact across teams via contracts.
3. 🟡 **Why automate governance instead of relying on review?** — Manual style/pattern review doesn't scale and drifts; lints/import-rules/CI enforce consistently and free reviews for substance.
4. 🟡 **What does a design system provide at scale?** — One shared UI vocabulary (tokens + components + theming) so features are consistent, aren't reinvented, and can be restyled globally in one place.

5. 🟡 **How does Conway's Law apply here?** — Module/team boundaries tend to mirror each other; align them deliberately (one team per module) so structure matches org.
6. 🔴 **How do you prevent pattern drift as the team grows?** — Automated enforcement (lints/boundaries/CI) + documented conventions/ADRs + a template example slice + required reviewers.
7. 🔴 **How do you right-size this?** — Light conventions + feature-first for small teams; modular packages + CODEOWNERS + strict governance + design system for large/multi-team — introduce as pressure appears.

Senior Engineer Tips

- Automate every convention you can (lints, import boundaries, format, CI); documented-but-unenforced conventions drift the moment the team grows past a handful.
- Align module ownership with team boundaries and route cross-team needs through contracts; that's how you get true parallel work instead of constant merge collisions.
- Invest early in a shared design system and an example slice as the canonical template; both pay back massively in consistency and onboarding as headcount rises.

Architect Perspective

Codebase and team scaling are two faces of the same problem, solved by the same structural moves: cohesive, boundaried, owned modules + automated governance + a shared design system. This deliberately aligns architecture with organization (Conway's Law), enabling many teams to work in parallel on a coherent codebase without collisions or drift — the human-scaling payoff of the feature-first/modular backbone. Right-sized to team size and enforced by tooling, it's what keeps a growing app buildable, consistent, and ownable (Module 44, Module 45, Module 50).

Summary

- Scale codebase + team together via feature-first/modular structure (cohesion + enforced boundaries), module ownership (CODEOWNERS), automated governance (lints/CI/docs/ADRs), and a shared design system.
- Goal: parallel work without collisions (boundaries + ownership) or drift (governance + design system); align module ≈ team (Conway's Law).
- Automate enforcement, provide a template slice + onboarding, and right-size to team size.

Revision Notes

- Structure: feature-first→modular + enforced boundaries; ownership: modules/packages ↔ teams (CODEOWNERS), contracts across teams; align module≈team (Conway).
- Governance: conventions + automated lints/import-rules/format/CI gates + ADRs + architecture doc + example-slice template; automate (not just docs).
- Design system: shared package (tokens + components + theming), central ownership, one UI vocabulary, global restyle in one place; right-size to team size.

Practice Questions

1. Which levers scale codebase and team together?
2. Why automate governance rather than rely on review?
3. What problems does a shared design system solve at scale?

Coding Questions

1. Write a CODEOWNERS mapping modules to teams (gated contracts).
2. Configure lints/import-boundary rules enforcing feature→core-only.
3. Sketch a design-system package API (tokens + a couple of components).

Mini Project

Codebase + team governance plan (Flutter): For a growing multi-team app, produce: a modular structure + CODEOWNERS mapping (modules↔teams, gated contracts/design-system), a governance stack (lints + import-boundary rules + CI gates + ADRs + example-slice template), and a shared design-system package outline (tokens + components + theming, central ownership). Right-size to team size. Acceptance: module/package ownership (aligned to teams); automated governance (not docs-only) + boundary enforcement; design system as shared UI vocabulary; onboarding/ADR/template noted; right-sizing to team scale; parallel-work-without-drift rationale.

Performance & Runtime Scaling

> As features, data, and users grow, the **runtime** dimension degrades in predictable ways: **startup** slows (more init work + bigger app), **memory** climbs (more state/caches/images/features resident), and the **frame budget** (~16ms at 60fps / ~8ms at 120fps) gets blown (bigger lists, more rebuilds, heavier screens). Scaling runtime means enforcing a **performance budget** (startup/memory/frame targets), **deferring work** (lazy init, deferred/route-based loading, on-demand feature loading), keeping the **frame budget** (virtualized lists, scoped rebuilds, const/repaint boundaries), and **profiling continuously** — so the app stays fast at 500 screens as it was at 5.

Introduction

This file covers keeping runtime healthy under growth: the performance budget, deferring startup/memory/loading work, protecting the frame budget at scale, and continuous profiling. It applies the performance techniques of Module 21 to the specific pressures of a *large, growing* app.

Why this concept exists

Individual optimizations aren't enough at scale — growth *systematically* erodes startup, memory, and frames unless you set budgets and defer work by default. Without a budget + deferral strategy, each new feature adds a little startup time and memory until the app is slow and bloated, with no single culprit. Runtime scaling is about **structural** performance, not one-off fixes.

Real-world analogy

A growing app is like a **car accumulating cargo**: add enough and acceleration (startup) suffers, fuel/space (memory) fills, and cornering (frame budget) degrades — gradually, invisibly. The fix is a **weight budget** (limits + weigh-ins = profiling), **only carrying what you need right now** (deferred/lazy loading), and **not loading the whole warehouse at the depot** (don't init every feature at startup). You monitor total weight continuously, not after it won't move.

Internal Working

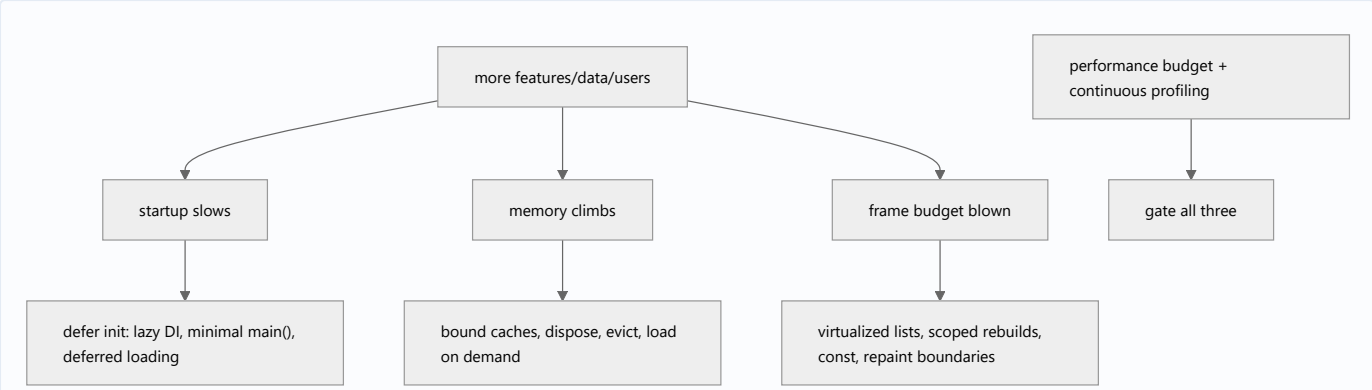

- **Set a performance budget** (targets you don't exceed): e.g., **cold start < X s**, **memory < Y MB** on target devices, **frame time ≤ 16ms (60fps)/8ms (120fps)** with jank < Z%, **app size < W MB**. Measure against it in CI/monitoring; a feature that blows the budget must be optimized/deferred (Module 21/Module 52).
- **Startup scaling** (init grows with features):
  - Keep `main()` **minimal**; **lazy-register DI** (create on first use, not at boot — Module 14); defer non-critical init (analytics, prefetch) to after first frame.
  - **Deferred/route-based loading**: don't build/load every feature upfront; load a feature's code/state **when navigated to** (lazy routes; Flutter web **deferred imports**/ `--split-debug-info` ; on-demand feature loading in modular apps).
  - Reduce **app size** (tree-shaking, `--obfuscate --split-debug-info` , split assets, app thinning — Module 21/Module 51).
- **Memory scaling** (more state/caches resident):
  - **Bound caches** (TTL/LRU/size — Module 34); **dispose** controllers/subscriptions/streams; **evict** off-screen feature state; use `cached_network_image` /resized images (image memory is a top offender).
  - **Load features/state on demand** and release when leaving; don't keep every feature's data resident.
  - Watch for **leaks** (undisposed listeners, retained contexts) — they compound at scale (Module 52).
- **Frame-budget scaling** (bigger/heavier UI):
  - **Virtualize long lists** (`ListView.builder` /slivers — never build all items); **scoped rebuilds** (select slices — Module 43); `const` **widgets** + `RepaintBoundary` ; avoid expensive work in `build` ; offload heavy compute to **isolates** (Module 02).
  - Keep per-screen widget trees shallow; avoid whole-page rebuilds on small changes.
- **Profile continuously** (not once): DevTools (timeline, memory, CPU), performance overlay, integration/perf tests in CI, and **production monitoring** (startup/frame/crash metrics —

Module 52). Catch regressions **as features land**, keyed to the budget.

- **Deferral by default (the scaling mindset)**: at scale, the default answer to "when to load/init X?" is **"as late as possible / on demand"** — because *every* feature adding a little upfront cost is what kills large apps.
- **Right-sizing**: a small app rarely needs budgets/deferred loading; a large one needs all of it. Introduce as the runtime dimension shows pressure (01_scaling_dimensions.md).

Memory Representation

The budget is a set of measured targets; deferral means feature code/state isn't resident until needed; bounded caches cap resident memory; disposal frees off-screen state. The invariant: resident work ≈ what's currently on screen/needed, not the whole app.

Compiler Behavior

Tree-shaking + deferred imports (web) + obfuscation/split-debug-info reduce shipped/loaded code; `const` enables widget canonicalization. Build tooling supports app-size + deferred-loading levers.

Runtime Behavior

Lazy init → faster cold start; on-demand loading → lower baseline memory + faster start; virtualized lists + scoped rebuilds + isolates → frames stay within budget; bounded caches + disposal → stable memory. Without these, all three degrade with growth.

Flutter Engine Behavior

Frame-budget levers reduce build/layout/paint/raster work per frame (Module 09); virtualization limits widget/element creation; repaint boundaries isolate repaints; the raster thread stays within budget.

Dart VM Behavior

Isolates offload heavy compute off the UI isolate; bounded memory reduces GC pressure; deferred loading defers code compilation/loading. Startup benefits from minimal boot-time work.

Examples

```
// Startup: minimal main() + lazy DI + defer non-critical work to after first frame
void main() {
  final di = GetIt.instance;
  registerCoreLazily(di);                 // lazy singletons: created on first use, not at boot
  runApp(const App());
}
// After first frame, kick off non-critical init (analytics/prefetch) — off the startup path.
WidgetsBinding.instance.addPostFrameCallback((_) => initNonCritical());

// Deferred/route-based feature loading (don't build every feature upfront)
GoRoute(path: '/reports', builder: (_, __) => const ReportsScreen()); // built only when navigated
// (Flutter web: deferred import of a heavy feature library, loaded on first use.)

// Frame budget: virtualized list + scoped rebuild + const
ListView.builder(                         // builds only visible items (virtualized)
  itemCount: items.length,
  itemBuilder: (_, i) => const _Row(),     // const where possible
);
Selector<FeedVm, int>(                     // scoped rebuild (Module 43)
  selector: (_, vm) => vm.unreadCount,
  builder: (_, count, __) => Badge(count),
);

// Memory: bounded cache + dispose
final imageCache = LruCache(maxBytes: 100 << 20);   // bound it
@override void dispose() { _sub.cancel(); _controller.dispose(); super.dispose(); } // release
```

Diagrams

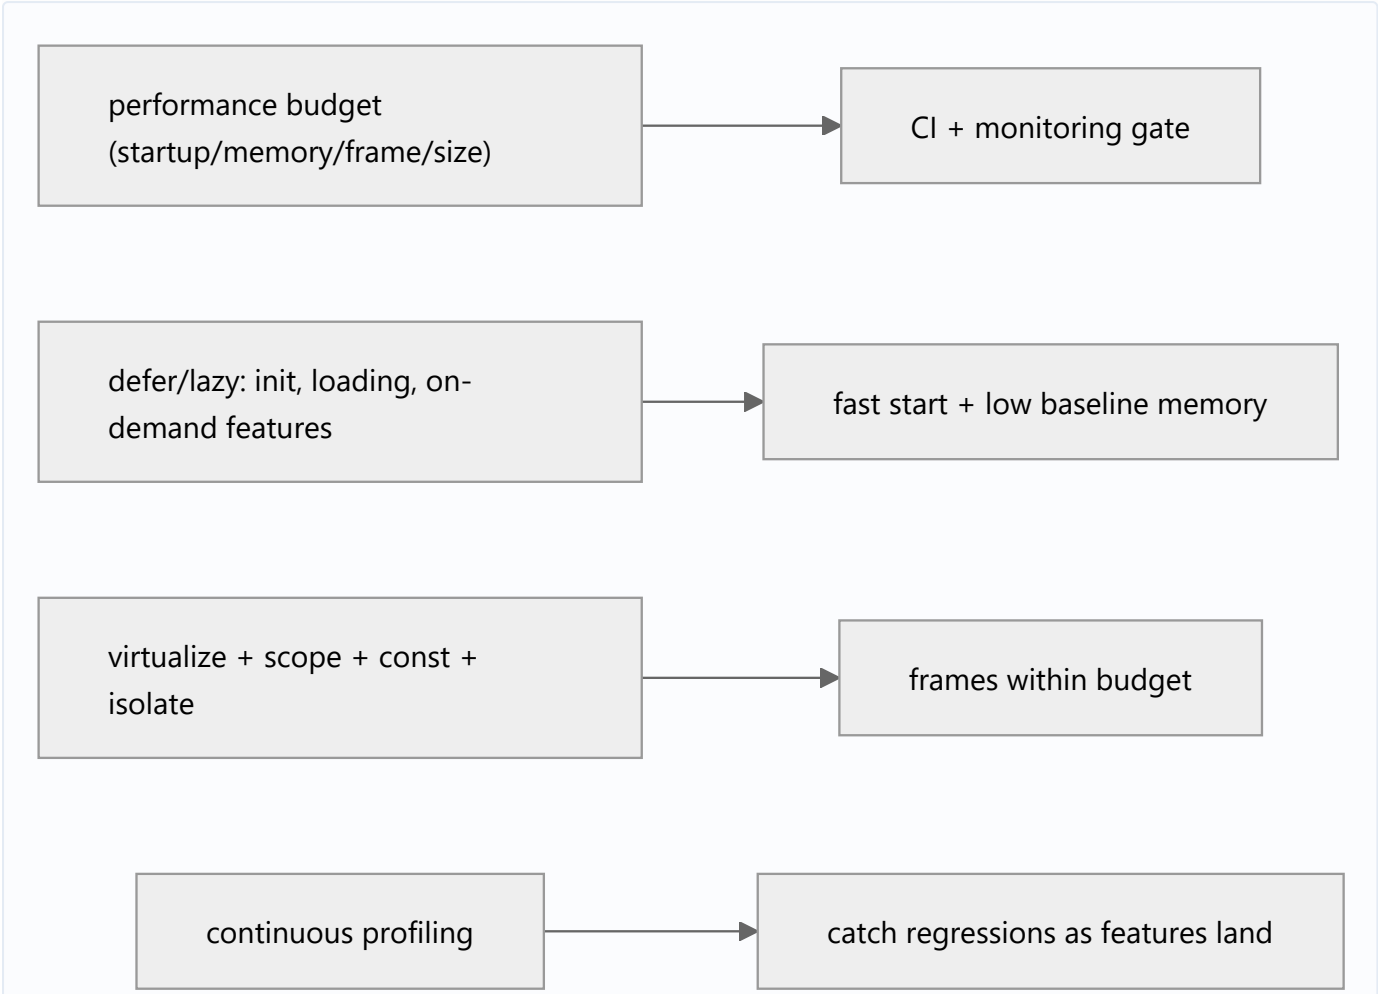

Common Mistakes

Mistake	Why it degrades at scale	Fix
Eager-initializing everything at boot	Slow cold start grows per feature	Lazy DI + defer non-critical + on-demand loading
Building all list items	Memory/jank on long lists	Virtualize (`.builder` /slivers)
Whole-page rebuilds	Frame budget blown	Scoped rebuilds + const + repaint boundaries
Unbounded caches / no disposal	Memory climbs + leaks	Bound caches (TTL/LRU) + dispose
Keeping all feature state resident	Baseline memory bloat	Load on demand; evict off-screen
Heavy compute in build/UI isolate	Jank	Offload to isolates
Profiling only once (or never)	Regressions slip in	Continuous profiling + CI/monitoring vs budget
No performance budget	No guardrail; gradual rot	Set + enforce startup/memory/frame/size targets

Best Practices

- Set and **enforce a performance budget** (cold start, memory, frame time/jank, app size) via **CI + production monitoring**; block features that blow it.
- **Defer by default**: minimal `main()`, **lazy DI**, post-first-frame non-critical init, **route-based/on-demand** feature + code loading; reduce app size.
- Protect the **frame budget**: **virtualize lists**, **scope rebuilds**, `const` + `RepaintBoundary`, offload heavy compute to **isolates**; keep resident work ≈ what's on screen.
- **Bound caches + dispose** everything (controllers/subscriptions/images); **profile continuously** and catch regressions as features land; **right-size** to the app's stage.

Performance

This whole file *is* the performance discipline for scale: budgets prevent gradual rot, deferral keeps startup/memory low as features multiply, frame levers keep 60/120fps under heavier UI, and continuous profiling catches regressions. The compounding win is **flat performance as the app grows**, instead of linear degradation.

Advantages / Disadvantages

- **+** Stable startup/memory/frames as features grow, regression detection, smaller/faster app, sustainable runtime scaling.
- **−** Budget + monitoring setup, deferral/laziness complexity (loading states, on-demand wiring), disposal discipline, profiling effort; overkill for tiny apps.

Interview Questions

1. 🟢 **How does growth degrade runtime, and along which axes?** — Startup slows (more init/bigger app), memory climbs (more resident state/caches), and the frame budget is blown (bigger lists/more rebuilds).
2. 🟢 **What is a performance budget?** — Enforced targets (cold start, memory, frame time/jank, app size) measured in CI + monitoring; features that exceed them must be optimized or deferred.
3. 🟡 **How do you keep startup fast as features multiply?** — Minimal `main()`, lazy DI, defer non-critical init to after first frame, and route-based/on-demand feature + code loading (deferred imports on web).
4. 🟡 **How do you protect the frame budget at scale?** — Virtualize lists, scope rebuilds, `const` / `RepaintBoundary`, and offload heavy compute to isolates — keep per-frame work small.
5. 🟡 **How do you keep memory bounded?** — Bound caches (TTL/LRU/size), dispose controllers/subscriptions, evict off-screen feature state, resize/cache images, and watch for

leaks.

6. 🔴 **Why is "defer by default" the scaling mindset?** — Because every feature adding a little upfront init/memory is what kills large apps; loading/initializing on demand keeps baseline cost flat.
7. 🔴 **Why profile continuously rather than once?** — Regressions accumulate as features land; continuous profiling (DevTools + CI perf tests + production monitoring) catches them against the budget.

Senior Engineer Tips

- Set a performance budget early and wire it into CI + monitoring; without a guardrail, large apps rot gradually with no single culprit to blame.
- Make deferral the default — lazy DI, on-demand feature loading, post-first-frame init; eager everything-at-boot is the number-one startup killer at scale.
- Treat unbounded caches, missing disposal, and un-virtualized lists as bugs; at scale they're the memory/jank offenders that compound silently.

Architect Perspective

Runtime scaling is structural performance: a budget defines the guardrails, deferral keeps baseline cost flat as features grow, frame/memory levers protect responsiveness under heavier UI/data, and continuous profiling enforces it all. Combined with modular on-demand loading and monitoring, it lets an app grow to hundreds of features while staying as fast as it was small. The architect's job is to institutionalize the budget + deferral defaults so performance is a property of the system, not a heroic one-off — proportional to the app's actual runtime pressure (Module 21, Module 52, 01_scaling_dimensions.md).

Summary

- Growth degrades startup, memory, and frames unless managed structurally: set + enforce a performance budget (CI + monitoring).
- Defer by default (lazy DI, on-demand feature/code loading, post-first-frame init); protect the frame budget (virtualize, scope, const, isolates); bound caches + dispose.
- Profile continuously to catch regressions as features land; right-size to the app's runtime pressure.

Revision Notes

- Budget: cold start / memory / frame time (16ms@60 / 8ms@120) + jank% / app size — enforce via CI + monitoring; block over-budget features.

- Startup: minimal `main()`, lazy DI, post-first-frame non-critical init, route-based/on-demand + deferred imports (web), reduce app size.
- Memory: bound caches (TTL/LRU), dispose subs/controllers, evict off-screen, resized/cached images, no leaks. Frame: virtualize lists, scope rebuilds, const + RepaintBoundary, isolates.
- Profile continuously (DevTools + CI perf + monitoring); defer-by-default mindset; right-size.

Practice Questions

1. What are the three runtime axes that degrade with growth, and their levers?
2. Why is "defer by default" essential at scale?
3. How do you enforce that a new feature doesn't blow the performance budget?

Coding Questions

1. Make `main()` minimal + lazy DI + post-first-frame non-critical init.
2. Convert an eager list + whole-page rebuild to virtualized + scoped.
3. Bound a cache + add disposal, and add a CI perf check against the budget.

Mini Project

Runtime scaling plan (Flutter): For a growing app, define a performance budget (cold start/memory/frame/app-size targets), implement deferral (minimal `main()` + lazy DI + post-first-frame init + route-based/on-demand feature loading), frame-budget levers (virtualized lists, scoped rebuilds, const/repaint boundaries, isolate for a heavy task), and memory hygiene (bounded caches + disposal), plus a continuous-profiling/CI-gate hook. Acceptance: enforced budget (CI + monitoring); deferral-by-default startup + loading; frame-budget levers applied; bounded caches + disposal; continuous profiling; right-sized to runtime pressure.

Technical Debt & Evolution

> Scaling is continuous change, and change accrues **technical debt** — shortcuts, outdated patterns, duplicated code, and deferred fixes that slow future work. Managing it is a discipline: **make debt visible** (a register/tracker, not tribal memory), **distinguish deliberate/prudent debt** (a conscious tradeoff you'll repay) from **reckless/inadvertent debt** (accidental rot), **pay it down continuously** (the boy-scout rule + scheduled refactoring, not a mythical "rewrite later"), and **evolve safely at scale** via **incremental refactoring, deprecation cycles, and strangler-fig migrations** — never big-bang rewrites. Debt isn't inherently bad; **unmanaged, invisible debt** is.

Introduction

This file covers the feature-scale + longevity dimension: tracking and classifying debt, paying it down, and evolving a large codebase safely (refactoring, deprecation, migration) — so growth doesn't calcify into un-changeable software. It complements the migration mechanics of Module 44/Module 45.

Why this concept exists

Every scaling app accumulates debt; the question is whether it's **managed** (visible, classified, repaid) or **ignored** (until velocity collapses and everyone fears touching the code). Unmanaged debt is the silent killer of large apps. A deliberate practice keeps the codebase changeable as it grows — turning "we can't touch that" into "we evolve it safely."

Real-world analogy

Tech debt is like **financial debt**: taken deliberately for leverage (ship now, refactor next sprint) it's a tool; ignored, the **interest compounds** until you can't afford new purchases (features). You keep a **ledger** (debt register), **make interest payments** (continuous refactoring), and **restructure large loans carefully** (planned migrations) — you don't declare bankruptcy and rebuild the whole company (big-bang rewrite), which usually fails.

Internal Working

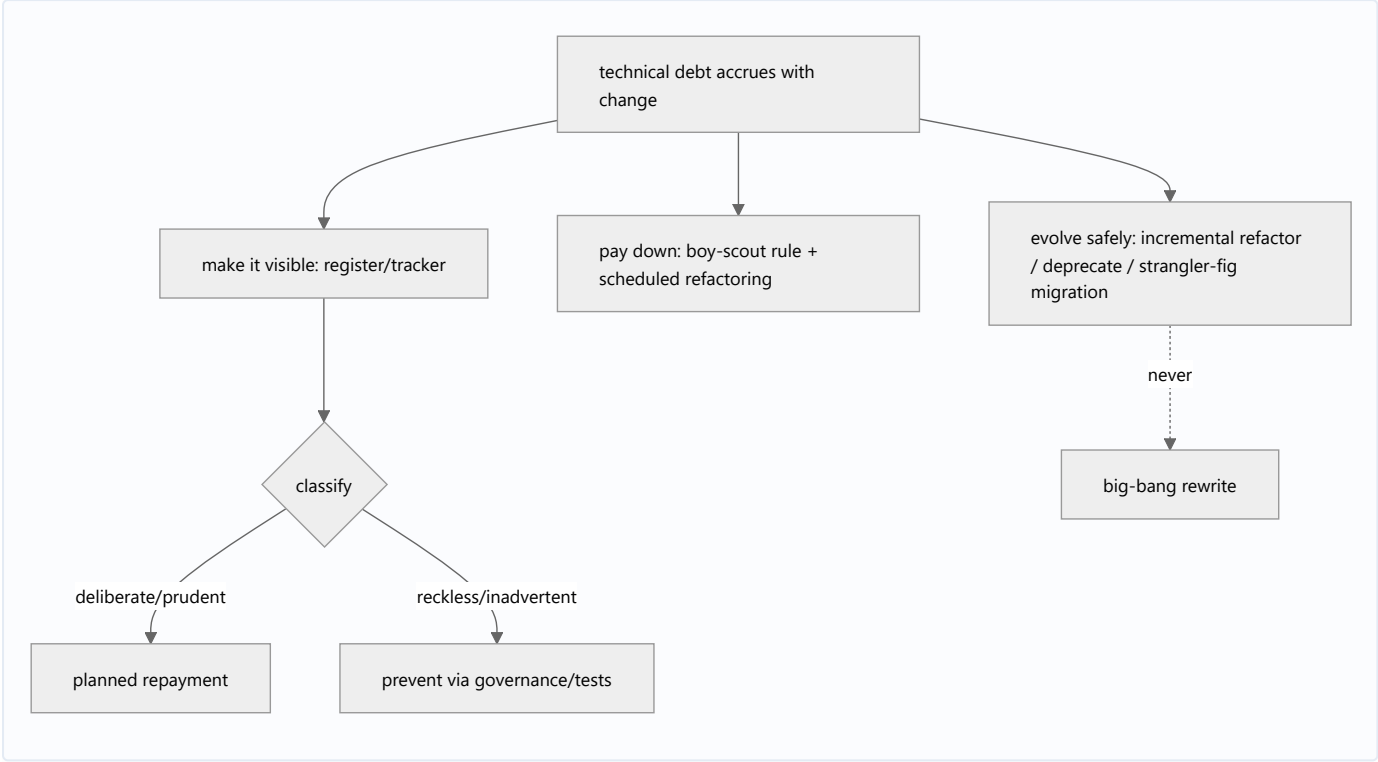

- **Make debt visible**: track it explicitly — a **debt register/backlog** (tickets tagged `tech-debt`), `// TODO` / `// DEBT` with links, ADRs recording deliberate shortcuts. Invisible debt (tribal knowledge) can't be prioritized or repaid.
- **Classify debt** (the quadrant — deliberate/inadvertent × prudent/reckless):
  - **Deliberate + prudent**: "we ship now with a known shortcut and a plan to repay" — legitimate leverage; record + schedule.
  - **Deliberate + reckless**: "no time for design" — dangerous; avoid.
  - **Inadvertent + prudent**: "now we know a better way" — learning; refactor as you go.
  - **Inadvertent + reckless**: incompetence/rot — prevent via governance/tests/review.
  - Knowing *which* kind guides response (repay planned vs prevent vs learn).
- **Pay it down continuously**:
  - **Boy-scout rule**: leave code a little better than you found it (small refactors alongside feature work).
  - **Scheduled capacity**: dedicate a % of each cycle to debt (not a mythical "big refactor later" that never comes).
  - **Tests first**: add tests around debt-laden code **before** refactoring so changes are safe (Module 49).
- **Refactoring safely (change structure, not behavior)**: small, behavior-preserving steps under test coverage; use tooling; keep PRs small; verify via tests/CI. Big untested refactors are the risk.

- **Deprecation cycles**: to evolve/replace an API, **deprecate** (mark `@Deprecated('use X')`, warn, document), migrate callers **incrementally**, then remove — never break consumers abruptly (esp. in modular apps where contracts have many consumers — Module 45).
- **Large migrations = strangler-fig, not big-bang**: replace old with new **incrementally** (new grows around old until old is gone), shipping throughout — the same pattern as layer→feature-first migration (Module 44). **Big-bang rewrites** are high-risk, long-freeze, and frequently fail — avoid them.
- **Prevent new debt**: governance (lints, review, architecture guide), tests, and design-system/contracts reduce the rate debt accrues (02_codebase_and_team_scaling.md).
- **Right-sizing**: don't chase zero debt (uneconomical) or ignore it (fatal); **manage** it — repay high-interest debt (in hot, frequently-changed areas), tolerate low-interest debt (stable, rarely-touched code).

Memory Representation

Not runtime — a **debt ledger** (tracked items, classified, prioritized by interest/impact) + an evolution process (deprecation timelines, migration plans). The invariant: debt is visible, classified, and being repaid faster than it accrues in hot areas.

Compiler Behavior

`@Deprecated` surfaces deprecation warnings guiding migration; lints/analysis prevent classes of new debt; tests + CI gate refactors.

Runtime Behavior

Not applicable — debt/evolution are dev-time concerns (though unmanaged debt indirectly causes runtime bugs/regressions and slows fixes).

Flutter Engine Behavior

Not applicable.

Dart VM Behavior

Not applicable.

Examples

```
// Making deliberate debt visible + planned
// DEBT(#1234, deliberate+prudent): hardcoded pricing; replace with PricingService by Q3.
// See ADR-017. Interest: HIGH (pricing changes weekly) -> prioritize repayment.

// Deprecation cycle: evolve an API without breaking consumers
@Deprecated('Use OrderRepository.findById. Removed in v3.0.')
Future<Order?> getOrder(String id) => findById(id);   // keep old, warn, migrate callers, then remove
```

```
Strangler-fig migration (incremental, ship throughout) — NOT big-bang:
  1) add the new implementation behind the same interface/contract
  2) route a slice of traffic/callers to the new path (flag/one feature)
  3) migrate callers incrementally; keep old working
  4) delete the old once no callers remain
  # coexist during transition; never a long code freeze

Refactor-safely loop:
  add tests around the debt -> small behavior-preserving change -> CI green -> repeat
```

Diagrams

Common Mistakes

Mistake	Why it's a problem	Fix
Invisible debt (tribal memory)	Can't prioritize/repay	Debt register + ADRs + tagged TODOs
"Big refactor later" (never happens)	Debt compounds	Continuous paydown (boy-scout + scheduled capacity)
Refactoring without tests	Behavior regressions	Add tests first; small steps
Big-bang rewrite	High risk, long freeze, often fails	Strangler-fig incremental migration
Breaking APIs abruptly	Breaks consumers (esp. contracts)	Deprecation cycle (mark → migrate → remove)
Chasing zero debt	Uneconomical	Manage by interest; tolerate low-interest debt
Ignoring debt in hot areas	Velocity collapse	Prioritize high-interest (frequently-changed) debt

Best Practices

- **Make debt visible** (register/tracker + ADRs + linked TODOs) and **classify** it (deliberate/prudent vs reckless/inadvertent) to guide response.
- **Pay down continuously** — boy-scout rule + scheduled capacity + **tests before refactoring**; prioritize **high-interest debt** in hot, frequently-changed areas.
- **Evolve safely at scale**: **deprecation cycles** for API changes (mark → migrate → remove), **incremental strangler-fig migrations** — **never big-bang rewrites**.
- **Prevent new debt** via governance/tests/design-system/contracts; **right-size** — manage debt (not zero, not ignored) by interest/impact.

Performance

Not runtime — the "performance" is **sustained development velocity**: managed debt keeps change cheap; unmanaged debt makes every change slow and risky. Prioritizing high-interest debt (hot areas) maximizes velocity ROI; refactoring stable low-interest code is low-value.

Advantages / Disadvantages

- **+** Sustained changeability + velocity, safe evolution, deliberate tradeoffs, no fear of touching code, avoided failed rewrites.
- **−** Requires discipline (tracking, scheduled capacity, tests), deprecation/migration overhead, judgment on interest/priority, ongoing effort.

Interview Questions

1. 🟢 **What is technical debt and is it always bad?** — Shortcuts/outdated patterns/deferred fixes that slow future work; not inherently bad — deliberate, prudent, managed debt is leverage; unmanaged/invisible debt is the problem.
2. 🟢 **How do you make debt manageable?** — Make it visible (register/ADRs/TODOs), classify it, and pay it down continuously (boy-scout rule + scheduled capacity) rather than a mythical big refactor.
3. 🟡 **How do you refactor safely at scale?** — Add tests around the code first, then make small behavior-preserving changes gated by CI — no big untested rewrites.
4. 🟡 **How do you change/replace an API without breaking consumers?** — A deprecation cycle: mark `@Deprecated` (with the replacement), migrate callers incrementally, then remove.
5. 🟡 **Why avoid big-bang rewrites?** — High risk, long code freeze, and frequent failure; use incremental strangler-fig migration and ship throughout.
6. 🔴 **How do you prioritize which debt to repay?** — By "interest": high-interest debt in hot, frequently-changed areas first; tolerate low-interest debt in stable code.

7. 🔴 **How do the debt quadrants guide response?** — Deliberate+prudent → schedule repayment; deliberate+reckless → avoid; inadvertent+prudent → refactor as you learn; inadvertent+reckless → prevent via governance/tests.

Senior Engineer Tips

- Track debt like tickets and repay it continuously (boy-scout rule + a fixed capacity slice); the "we'll do a big cleanup later" plan is how codebases calcify.
- Always add tests before refactoring and change behavior in small, CI-gated steps; untested refactors are how "cleanup" becomes an outage.
- Evolve via deprecation + strangler-fig and prioritize by interest (hot areas); big-bang rewrites and abrupt API breaks are the two most expensive mistakes at scale.

Architect Perspective

Managing technical debt and enabling safe evolution is what keeps a scaling app *changeable* over years: visible, classified debt repaid continuously by interest, APIs evolved through deprecation, and large changes made via incremental strangler-fig migration rather than doomed rewrites. It's the temporal dimension of scaling — complementing structure (codebase/team) and runtime — and it's sustained by the same governance/tests/contracts that prevent debt from accruing. The architect institutionalizes the ledger + paydown + evolution process so growth compounds capability instead of calcifying it (02_codebase_and_team_scaling.md, Module 44, Module 49).

Summary

- Tech debt is inevitable with change; manage it — make it visible (register/ADRs), classify it, and pay it down continuously (boy-scout + scheduled capacity, tests first).
- Evolve safely: deprecation cycles for APIs, incremental strangler-fig migrations — never big-bang rewrites.
- Prioritize by interest (hot areas), prevent new debt via governance/tests/contracts; manage (not zero, not ignored).

Revision Notes

- Make debt visible (register/tracker + ADRs + linked TODOs); classify (deliberate/prudent vs reckless/inadvertent → guides response); not inherently bad.
- Pay down continuously (boy-scout rule + scheduled capacity + tests-before-refactor); prioritize high-interest debt (hot/frequently-changed areas); tolerate low-interest.

- Evolve: `@Deprecated` cycles (mark→migrate→remove), strangler-fig incremental migration (ship throughout), never big-bang; prevent new debt via governance/tests/contracts.

Practice Questions

1. When is technical debt acceptable, and when is it dangerous?
2. How do you refactor a debt-laden module safely?
3. Why deprecation + strangler-fig over a big-bang rewrite?

Coding Questions

1. Record a deliberate debt item (register/ADR/TODO) with interest + repayment plan.
2. Add a deprecation cycle to evolve an API without breaking callers.
3. Outline a strangler-fig migration (with test-first refactoring) for a legacy module.

Mini Project

Debt + evolution plan (Flutter): For a scaling app, create a technical-debt register (a few items, classified by quadrant + interest, with repayment plans), a continuous-paydown policy (boy-scout rule + scheduled capacity + tests-first), a deprecation-cycle example for an evolving API, and a strangler-fig migration plan for a legacy module (incremental, ship-throughout, no big-bang). Acceptance: debt made visible + classified + prioritized by interest; continuous paydown policy (not "later"); safe refactoring (tests-first, small steps); deprecation cycle; strangler-fig migration (no big-bang); prevention via governance/tests noted.

Scalable Integration (Capstone: A Scaling Playbook)

> Synthesize the four dimensions into one **stage-aware scaling playbook**: for a given app stage, it prescribes the **structure** (feature-first→modular), **team practices** (ownership, governance, design system), **runtime discipline** (performance budget + deferral), and **evolution strategy** (debt register + deprecation + strangler-fig) — each **right-sized**, with explicit **"do now / defer"** calls and the **symptoms that trigger the next investment**. The playbook is the deliverable: not maximal architecture, but a diagnosed, proportional plan that keeps codebase, team, runtime, and features healthy as the app grows.

Introduction

This module capstone assembles the four dimensions + their levers into a practical, stage-aware playbook — the artifact a lead uses to decide what to invest in now versus later. It ties together the whole architecture band as a "how to grow deliberately" guide.

Why this concept exists

Scaling knowledge is only useful applied *proportionally* to a real app's stage. A playbook converts the dimensions/levers into concrete, staged decisions (with triggers), preventing both over-engineering (max architecture on a small app) and reactive fire-fighting (fixing a dimension only when it's already on fire).

Real-world analogy

It's a **city master plan with phases**: phase 1 (village) needs basic roads + a shared water source; phase 2 (town) adds districts + zoning + utilities capacity; phase 3 (city) adds transit, governance, and redevelopment plans for aging areas. Each phase is **triggered by measured growth**, invests only what's needed now, and plans the next phase's triggers — not building a metro system for a village.

Internal Working

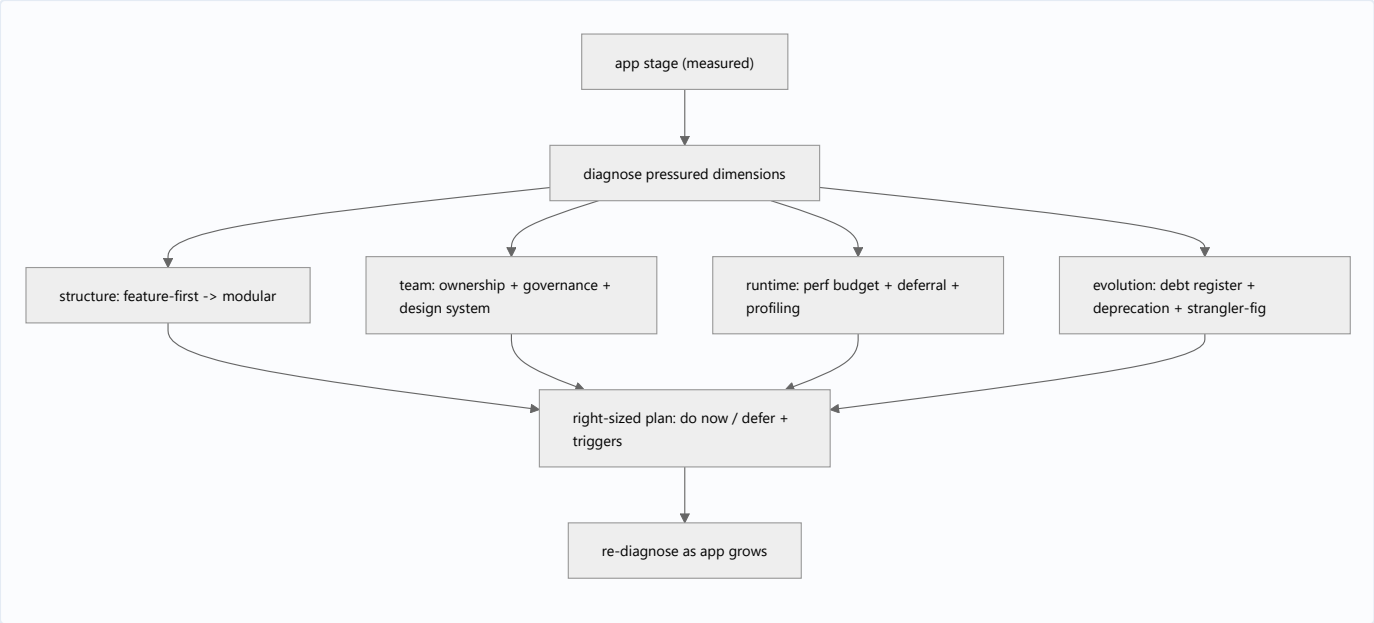

- **Stage-aware, right-sized** (01_scaling_dimensions.md): the playbook adapts to stage — **prototype** (minimal), **growing single-team**, **large multi-team** — applying the **minimum lever** per pressured dimension and **deferring** the rest, with the **symptoms that trigger** the next step.
- **Structure lever** (codebase): start **feature-first** (cohesive slices + `core`); graduate to **modular packages** (enforced boundaries + incremental builds) when build/team pressure appears (Module 44/Module 45); use **Clean/MVVM** inside slices (Module 40/Module 43); apply **DDD** to the complex core only (Module 46).
- **Team lever**: **ownership** (CODEOWNERS/packages), **governance** (conventions + lints + import boundaries + CI gates + ADRs + template slice), a **shared design system** — introduced as contributor count grows (02_codebase_and_team_scaling.md).
- **Runtime lever**: a **performance budget** (startup/memory/frame/app-size) enforced by **CI + monitoring**, **deferral-by-default** (lazy DI, on-demand feature/code loading), frame-budget levers (virtualize/scope/const/isolate), bounded caches + disposal, **continuous profiling** (03_performance_and_runtime_scaling.md/Module 21/Module 52).
- **Evolution lever** (feature/longevity): a **debt register** (classified, interest-prioritized), **continuous paydown** (boy-scout + scheduled capacity + tests-first), **deprecation cycles**, **strangler-fig migrations** — never big-bang (04_technical_debt_and_evolution.md).
- **Cross-cutting hooks**: **CI/CD** (incremental, gated — Module 50), **monitoring** (perf/crash/usage — Module 52), **testing** (regression safety at scale — Module 49) — these enforce the playbook.
- **Do-now / defer + triggers**: for each dimension, the playbook states **what to do at this stage**, **what to defer**, and the **measured symptom** that promotes a deferred item to "do now" — making scaling **proactive-but-proportional**, not reactive.

- **Re-diagnose**: the bottleneck moves as the app grows; the playbook is **revisited** periodically against fresh measurements (01_scaling_dimensions.md).

Memory Representation

Not runtime — a **living document**: per-dimension current-state, chosen levers (do-now), deferred items + triggers, and cross-cutting hooks. Revisited as measurements change.

Compiler Behavior

Structure/governance levers are compile/build-enforced (modular boundaries, lints); the rest are process/runtime/monitoring — together they make much of the playbook automatically enforced.

Runtime Behavior

Only the runtime lever affects runtime (budget/deferral/profiling); the others shape build/dev/team/evolution. Monitoring feeds runtime + evolution decisions.

Flutter Engine Behavior

Runtime levers interact with the engine (rebuilds/rasterization/memory — Module 21); others don't.

Dart VM Behavior

Modular structure enables incremental builds; runtime levers touch isolates/GC; the rest are org/process.

Examples

```
Scaling playbook (excerpt) — do-now / defer / trigger, per stage:

STAGE: growing single-team app
  Structure   DO NOW: feature-first + core + Clean/MVVM inside slices
              DEFER: modular packages          TRIGGER: builds >5min OR 2nd team joins
  Team        DO NOW: conventions + lints + CI gates + a lightweight design system
              DEFER: strict CODEOWNERS/packages TRIGGER: >~8 contributors / merge-conflict pain
  Runtime     DO NOW: perf budget + deferral defaults + continuous profiling
              DEFER: on-demand code loading (web deferred imports) TRIGGER: cold start > budget
  Evolution   DO NOW: debt register + boy-scout + tests-first refactoring
              DEFER: formal deprecation policy TRIGGER: first public/shared contract
  Cross-cut   DO NOW: CI (analyze/test) + basic monitoring (crash/perf)

STAGE: large multi-team app -> promote most "defer" items to "do now" (modular packages,
       CODEOWNERS, design-system package, on-demand loading, deprecation cycles, incremental CI).
```

```
Reassessment loop:
  measure (build time, conflict rate, frame/startup/memory, regression rate)
  -> re-diagnose pressured dimension -> promote a deferred lever -> update playbook
```

Diagrams

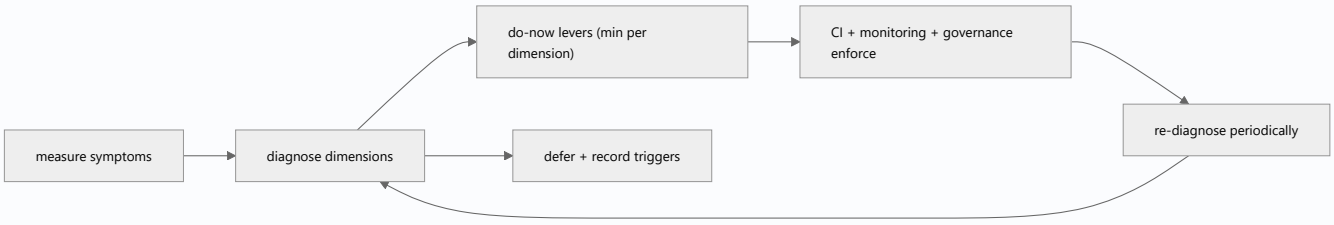

Common Mistakes

Mistake	Why	Fix
One-size playbook (max architecture)	Over-engineers small apps	Stage-aware, right-sized do-now/defer
No triggers for deferred items	Reactive fire-fighting later	Define measured symptoms that promote items
Investing in a non-bottleneck dimension	Wasted effort	Diagnose from symptoms first
Playbook as a one-time doc	Bottleneck moves	Living doc; re-diagnose periodically
No enforcement hooks	Levers erode	Wire CI/monitoring/governance
Ignoring evolution/debt	Calcifies at scale	Include debt register + migration strategy
Skipping measurement	Dogma over data	Measure build/conflict/frame/regression

Best Practices

- Produce a **stage-aware, right-sized playbook**: per dimension, state **do-now** levers, **deferred** items, and the **measured trigger** that promotes each — proactive but proportional.
- Cover all four dimensions with the band's levers (**feature-first→modular + Clean/MVVM/DDD**; **ownership + governance + design system**; **perf budget + deferral + profiling**; **debt register + deprecation + strangler-fig**).
- **Enforce** via cross-cutting hooks (**CI, monitoring, testing, governance**); **measure** symptoms and **re-diagnose** as the app grows.
- **Never over-engineer** (max architecture on a small app) or **fire-fight** (fix only when on fire); manage the tradeoffs continuously.

Performance

The playbook itself optimizes **investment ROI**: effort goes to the pressured dimension, deferred elsewhere, so the app scales without wasted work or gradual rot. Runtime performance is one dimension it governs (budget/deferral); the others govern dev/team/evolution velocity.

Advantages / Disadvantages

- **+** Proportional, diagnosed, proactive-but-not-wasteful scaling across all four dimensions; clear triggers; enforced + revisited; avoids over/under-engineering.
- **−** Requires measurement + judgment + upkeep (living doc); no formulaic answer; needs org buy-in for governance/ownership/budget enforcement.

Interview Questions

1. 🟢 **What is a scaling playbook?** — A stage-aware, right-sized plan covering all four scaling dimensions with do-now levers, deferred items, and measured triggers to promote them.
2. 🟢 **Why stage-aware/right-sized rather than one-size?** — To avoid over-engineering small apps and fire-fighting large ones; invest the minimum in the pressured dimension and defer the rest.
3. 🟡 **Which levers cover the four dimensions?** — Structure (feature-first→modular + Clean/MVVM/DDD), team (ownership + governance + design system), runtime (perf budget + deferral + profiling), evolution (debt register + deprecation + strangler-fig).
4. 🟡 **What promotes a deferred item to "do now"?** — A measured symptom/trigger (e.g., builds >5min → modularize; cold start > budget → on-demand loading; >N contributors → CODEOWNERS/design-system).
5. 🟡 **How is the playbook enforced?** — Cross-cutting hooks: CI (incremental/gated), monitoring (perf/crash/usage), testing (regression safety), and automated governance.
6. 🔴 **Why is the playbook a living document?** — The bottleneck moves as the app grows; you re-measure and re-diagnose periodically, promoting deferred levers as triggers fire.
7. 🔴 **How does the playbook synthesize the architecture band?** — It sequences Clean/MVVM/feature-first/modular/DDD + performance/testing/CI/monitoring as staged, dimension-targeted investments — proportional application of the whole band.

Senior Engineer Tips

- Write the playbook as do-now/defer/trigger per dimension and keep it living; the triggers are what turn scaling from reactive panic into scheduled, proportional investment.
- Diagnose from measurements (build time, conflict rate, frame/startup/memory, regression rate) before investing; most "we need to re-architect everything" is one dimension away from a targeted fix.
- Enforce every lever with a hook (CI, monitoring, governance) — an un-enforced playbook decays into aspiration; and revisit it each time the app crosses a growth threshold.

Architect Perspective

The scaling playbook is the architect's synthesis of the entire band: a stage-aware, measured, proportional plan that invests in the pressured dimension, defers the rest with explicit triggers, and enforces itself through CI/monitoring/governance. It reframes "architecture" from a fixed choice to a **continuous, diagnosed practice** — applying Clean/MVVM/feature-first/modular/DDD + performance/testing/evolution exactly when and where they pay off. Done

well, an app grows to hundreds of features and many teams while staying navigable, ownable, fast, and changeable — the definition of a scalable application (01_scaling_dimensions.md, Module 45, Module 50, Module 52).

Summary

- A scaling playbook is a stage-aware, right-sized plan across the four dimensions, with do-now levers, deferred items, and measured triggers.
- Levers: feature-first→modular + Clean/MVVM/DDD; ownership + governance + design system; perf budget + deferral + profiling; debt register + deprecation + strangler-fig.
- Enforce via CI/monitoring/testing/governance; measure + re-diagnose continuously; avoid over-engineering and fire-fighting.

Revision Notes

- Playbook = stage-aware, right-sized, per-dimension do-now/defer/trigger (proactive-but-proportional; living doc).
- Levers: structure (feature-first→modular + Clean/MVVM/DDD-core-only), team (ownership/governance/design system), runtime (perf budget + deferral + profiling), evolution (debt register + deprecation + strangler-fig).
- Enforce via CI/monitoring/testing/governance; measure symptoms (build/conflict/frame/startup/memory/regression) → re-diagnose → promote deferred levers; never over-engineer or fire-fight.

Practice Questions

1. What does a scaling playbook contain, and why do-now/defer/trigger?
2. Which levers map to which dimensions?
3. How is the playbook enforced and kept current?

Coding Questions

1. Draft a do-now/defer/trigger playbook for a given app stage.
2. Map measured symptoms to the levers they should promote.
3. List the CI/monitoring/governance hooks that enforce the playbook.

Mini Project

Scaling playbook (capstone — Flutter): For a chosen app stage (e.g., growing single-team heading toward multi-team), produce a scaling playbook covering all four dimensions: structure (feature-first now, modular deferred + trigger), team (governance/design-system now, CODEOWNERS/packages deferred + trigger), runtime (perf budget + deferral + profiling now, on-demand code loading deferred + trigger), and evolution (debt register + boy-scout/tests-first now, formal deprecation deferred + trigger) — plus CI/monitoring/testing hooks and a reassessment loop. Acceptance: all four dimensions covered with do-now/defer/measured-trigger; right-sized to the stage (no over-engineering); enforcement hooks (CI/monitoring/governance); living-doc reassessment loop; synthesizes the architecture band proportionally.