



Zero Downtime Migration Strategies for Legacy Angular Applications to Next Generation Framework Versions

Harshika Juvvadi

Full stack Developer, USA

Publication History: Received: 27.06.2026; Revised: 5.08.2026; Accepted: 10.08.2026; Published: 22.08.2026.

ABSTRACT: A production Angular application that has accumulated years of active use rarely has the luxury of a maintenance window long enough for a full rewrite, which leaves a team migrating that application toward a next generation framework version, whether from an early AngularJS codebase or from an older Angular major version, needing a migration strategy that keeps the application serving live traffic throughout the transition rather than treating migration as a discrete event. This article presents a zero downtime migration pipeline for legacy Angular applications, built around four coordinated elements, an incremental, route level cutover following the strangler fig pattern rather than a single big bang rewrite, dual version coexistence through a hybrid bootstrap that lets legacy and migrated code run side by side during the transition, automated codemods and schematics that mechanically apply structural changes such as standalone component conversion across many files consistently, and a feature flagged, canary gated rollout that exposes a newly migrated route to a small share of production traffic before a full cutover, with automated rollback if health checks fail. Each element is presented with illustrative workflow diagrams and code, including a hybrid bootstrap configuration, an illustrative codemod script, a feature flag and canary routing configuration, and a continuous integration configuration enforcing automated rollback. An illustrative case study for a single production Angular application reports simulated indicators showing the share of application routes fully migrated rising from an illustrative four percent under planning alone to ninety eight percent as the pipeline's four elements are adopted cumulatively, and user facing incidents during cutover windows falling from a mid single digit weekly count to near zero within roughly one and a half migration weeks of adoption. The article closes with a discussion of what the pipeline does not eliminate, the cost of the pipeline at small scale, limitations of the illustrative evaluation, and directions for further work. The intended audience is frontend engineers, platform and migration leads, and researchers studying legacy application modernization for Angular teams in 2026.

KEYWORDS: Angular, legacy migration, strangler fig pattern, feature flags, canary release, zero downtime, standalone components, continuous delivery

I. INTRODUCTION

An Angular application still serving real production traffic after several years of active development accumulates two kinds of debt simultaneously, the ordinary kind any codebase accumulates, and a specific migration debt tied to whichever framework version or architectural pattern was current when a given piece of the application was first written. A route built when AngularJS was the only option, or a module built before standalone components and signals existed, is not simply older code, it is code written against assumptions the framework itself has since moved past, and the gap between those assumptions and the framework's current direction only widens the longer migration is deferred.

This article treats migrating such an application as a continuous delivery problem rather than a one time event, on the premise that a production application's availability requirement does not pause for the migration's duration, which means the migration strategy itself must be designed around keeping the application serving traffic throughout rather than around minimizing the migration's own elapsed time at the expense of availability. Section 4 presents a four part pipeline built on that premise, moving from route level strangler fig cutover through hybrid version coexistence, automated codemods, and a feature flagged, canary gated rollout.

Two specific consequences of migrating without this structure motivate the pipeline's design. First, a big bang rewrite that replaces the entire application at once concentrates all of the migration's risk into a single cutover moment, which



means any defect introduced during the rewrite affects the entire application's availability simultaneously rather than being contained to the specific route or module where it originated, a problem Section 4.2's incremental, route level cutover addresses directly. Second, a migration with no automated rollback mechanism depends entirely on a human noticing a problem quickly enough to intervene manually, which is slower and less reliable than an automated health check acting the moment a defined threshold is crossed, a problem Section 4.6's canary gate addresses.

The remainder of the article is organized as follows. Section 2 reviews background on the strangler fig pattern, continuous delivery and canary releasing, feature flags, Angular's own migration tooling and version history, and semantic versioning as it applies to a phased migration. Section 3 states the specific problems the pipeline addresses. Section 4 presents the pipeline in seven parts with reference diagrams and illustrative code. Section 5 walks through an illustrative case study for one production Angular application with simulated indicators. Section 6 discusses what the pipeline does not eliminate and the cost of the pipeline at small scale, Section 7 discusses limitations, Section 8 suggests future work, and Section 9 concludes.

II. BACKGROUND AND LITERATURE REVIEW

2.1 The Strangler Fig Pattern

Fowler's description of the strangler fig pattern, named for a fig species that grows around a host tree and gradually replaces it without the host ever being removed all at once, gives this article's Section 4.2 its organizing idea, that a legacy system can be replaced incrementally, one bounded piece at a time, with the legacy and replacement systems coexisting throughout the transition, rather than requiring the legacy system to be fully replaced before any part of the replacement can go live (Fowler, 2004). Applied to an Angular application specifically, the pattern suggests migrating one route or module at a time while the surrounding application continues running on its existing version, which is precisely the incremental unit Section 4.2's cutover strategy adopts.

2.2 Continuous Delivery and Deployment Risk

Humble and Farley's account of continuous delivery argues that deployment risk is best managed by making deployments small, frequent, and automatically verifiable, rather than by making them infrequent and therefore individually larger and riskier, a principle this article's entire pipeline structure follows by decomposing a single large migration into many small, independently verifiable route level cutovers rather than one large migration event (Humble and Farley, 2010).

2.3 Canary Releasing

Google's site reliability engineering practice describes canary releasing as exposing a change to a small subset of production traffic first, monitoring defined health signals, and only proceeding to a full rollout if those signals stay within acceptable bounds, a pattern this article's Section 4.6 applies specifically to a migrated route's initial exposure, treating a migrated route the same way a canary release treats any other risky production change (Beyer et al., 2016).

2.4 Feature Flags for Controlled Rollout

Hodgson's account of feature flags, technique for shipping code to production behind a runtime toggle that controls whether a given user or request path exercises the new code, gives this article's Section 4.4 the mechanism for exposing a migrated route to a controlled share of traffic without requiring a separate deployment for each rollout percentage change, decoupling the decision of how much traffic sees the migrated route from the act of deploying the migrated code itself (Hodgson, 2017).

2.5 AngularJS End of Life and the Case for Migration

The Angular team's own announcement discontinuing long term support for AngularJS at the end of 2021 removed continued security patching as a reason to delay migration for any AngularJS application still in production, giving this article's motivating scenario in Section 1 a concrete deadline rather than only an abstract technical debt argument (Angular, 2022). For an AngularJS application specifically, Angular's own upgrade guide documents the ngUpgrade hybrid bootstrap mechanism this article's Section 4.3 builds on, allowing AngularJS and Angular code to run within the same application simultaneously during a phased migration (Angular, 2023).

2.6 Angular's Own Update Tooling and Recent Architectural Shifts

Angular's official update guide and its ng update command apply many version to version migration steps automatically through schematics, code transformation scripts that rewrite a project's source to match a new version's conventions, which Section 4.5's automated codemod step extends conceptually to structural migrations, such as standalone component conversion, that go beyond what the framework's own built in schematics cover (Angular, 2024). The introduction of



standalone components as a stable feature in Angular 15 removed the NgModule as a mandatory organizing unit for a component, giving a team migrating an older, NgModule based application a concrete target architecture to migrate toward (Gechev, 2022a). The earlier introduction of the Ivy rendering and compilation pipeline in Angular 9 similarly changed the framework's internal architecture substantially while preserving most application level backward compatibility, demonstrating that a major internal architectural shift and application level stability are not mutually exclusive goals (Raible, 2020).

2.7 Semantic Versioning for Phased Migration Tracking

Semantic versioning's convention of communicating breaking changes through a major version increment gives a phased migration a way to track which of an application's dependencies have been upgraded to a version compatible with the target architecture and which remain on an older, incompatible major version, a tracking need Section 4.7's telemetry and verification step depends on directly (Preston-Werner, 2013).

Table 1. Summary of prior work and its relationship to this article's pipeline

Prior work	Core contribution	Relationship to this pipeline
Fowler (2004)	Incremental legacy replacement with coexisting systems	Basis for the route level cutover strategy (Section 4.2)
Humble and Farley (2010)	Small, frequent, automatically verifiable deployments	Basis for decomposing migration into small cutovers
Beyer et al. (2016)	Canary releasing with automated health signal monitoring	Basis for the canary gate and automated rollback (Section 4.6)
Hodgson (2017)	Feature flags decoupling rollout from deployment	Basis for the feature flagged rollout mechanism (Section 4.4)
Angular (2022); Angular (2023)	AngularJS end of life and the ngUpgrade hybrid bootstrap	Basis for dual version coexistence (Section 4.3)
Angular (2024); Gechev (2022a)	Automated update schematics and standalone components	Basis for automated codemods and the migration target (Section 4.5)
Raible (2020)	Major internal architectural change with preserved compatibility	Precedent for the Ivy rendering pipeline transition
Preston-Werner (2013)	Semantic versioning convention for tracking breaking changes	Basis for dependency migration tracking (Section 4.7)

III. PROBLEM STATEMENT

The pipeline proposed in Section 4 responds to four specific, recurring problems observed in teams migrating a production Angular application without deliberate zero downtime structure. Each problem is stated here in general terms and revisited concretely in the illustrative case study in Section 5.

3.1 Big Bang Rewrites Concentrating Risk into a Single Cutover

A migration strategy that replaces the entire application at once, however carefully tested beforehand, concentrates every category of migration risk into a single cutover moment, which means a defect that testing missed affects the application's entire user base simultaneously rather than being contained to a specific, bounded piece of functionality.

3.2 No Mechanism for Legacy and Migrated Code to Coexist

Without an explicit coexistence mechanism, a team is forced to choose between migrating a piece of functionality completely before any of it can go live, which delays value delivery until the entire piece is done, or leaving the application in an inconsistent state where some functionality has been migrated and some has not with no defined boundary between them, which risks unpredictable interaction between the two.



3.3 Manual, Inconsistent Application of Structural Migration Changes

A structural change such as converting a component from an NgModule based declaration to a standalone declaration is mechanical in principle but tedious and error prone to apply by hand across a codebase with many components, and manual application at scale tends to produce inconsistent results, some components converted correctly, others converted with a subtle deviation from the intended pattern, which itself becomes a new source of technical debt.

3.4 No Automated Signal for When a Migrated Route Is Unsafe

A team that exposes a migrated route to production traffic with no automated health monitoring depends on a human noticing a problem, whether through a manual dashboard check or a user complaint, before deciding to roll back, which is slower than an automated system reacting the moment a defined health threshold is crossed and risks a longer window of degraded service than an automated gate would allow.

Table 2. Summary of the four pipeline problems and where Section 4 addresses each

Problem	Root cause	Addressed in
Big bang rewrites concentrate risk	Entire application replaced in a single cutover	Section 4.2, strangler fig route level cutover
No coexistence mechanism	No defined boundary between legacy and migrated code	Section 4.3, dual version coexistence with hybrid bootstrap
Manual structural changes are inconsistent	No automated, repeatable transformation mechanism	Section 4.5, automated codemods and schematics
No automated unsafe route signal	No health monitoring tied to an automated rollback action	Section 4.6, canary release and automated rollback gate

IV. PROPOSED ZERO DOWNTIME MIGRATION PIPELINE

4.1 Pipeline Principles

The pipeline rests on four principles that map directly onto the four problems in Section 3. First, migration proceeds route by route or module by module, following the strangler fig pattern, so risk stays bounded to whichever piece is currently being migrated. Second, legacy and migrated code coexist within the same running application throughout the migration, through an explicit hybrid bootstrap mechanism rather than an ad hoc or undefined boundary. Third, structural changes that apply consistently across many files are automated through codemods and schematics rather than applied by hand. Fourth, every migrated route is exposed to production traffic gradually, behind a feature flag and a canary health gate, with automated rollback if the gate's health checks fail.

Figure 1. Zero Downtime Migration Pipeline Overview




Figure 1. Zero downtime migration pipeline overview, from the legacy application through parallel migration mechanisms to a feature flagged, canary gated rollout and full cutover.




4.2 Strangler Fig Route Level Cutover

Rather than migrating the entire application at once, the pipeline identifies a bounded unit, typically a single route or a cohesive feature module, as the current migration target, following the incremental replacement strategy discussed in Section 2.1, and completes that unit's migration, verification, and full cutover before beginning the next one. Listing 1 shows an illustrative routing configuration directing traffic to either the legacy or the migrated implementation of a given route.

```
// app.routes.ts (illustrative excerpt)
export const routes: Routes = [
  {
    path: 'account-settings',
    loadComponent: () =>
      migrationFlags.isMigrated('account-settings')
        ? import('./account-settings/account-settings.component')
          .then((m) => m.AccountSettingsComponent)
        : import('./legacy/account-settings/upgraded.component')
          .then((m) => m.LegacyAccountSettingsComponent),
  },
];
```

Listing 1. An illustrative routing configuration selecting between the legacy and migrated implementation of a route based on migration flag state (Fowler, 2004).

Figure 2. Migration Stages from Route Identification to Cutover



A canary that fails its health checks at step four rolls back to the legacy route automatically, which the hybrid bootstrap from step two keeps available throughout the rollout (Section 4.6).

Figure 2. Migration stages from route identification to cutover, showing how a hybrid bootstrap, an automated codemod, and a feature flagged canary rollout combine to retire a legacy route.

4.3 Dual Version Coexistence with a Hybrid Bootstrap

For an application migrating from AngularJS specifically, the pipeline uses the ngUpgrade hybrid bootstrap mechanism discussed in Section 2.5 to run AngularJS and Angular within the same application simultaneously, with a defined set of upgraded and downgraded components bridging between the two, so that a route can be migrated independently of every other route rather than requiring the entire application to be on one framework version at a time. Listing 2 shows an illustrative hybrid bootstrap configuration.

```
// hybrid-bootstrap.ts (illustrative excerpt)
import { UpgradeModule } from '@angular/upgrade/static';
import { platformBrowser } from '@angular/platform-browser';
import { AppModule } from './app.module';

platformBrowser()
  .bootstrapModule(AppModule)
  .then((platformRef) => {
    const upgrade = platformRef.injector.get(UpgradeModule);
    upgrade.bootstrap(document.body, ['legacyAngularJsApp']);
  });
```



Listing 2. An illustrative ngUpgrade hybrid bootstrap, running the legacy AngularJS application and the Angular application together during migration (Angular, 2023).

4.4 Feature Flagged Rollout of Migrated Routes

Once a route's migration is complete and passing its own tests, the pipeline uses a feature flag, following the decoupling principle discussed in Section 2.4, to control what share of production traffic reaches the migrated implementation rather than the legacy one, allowing that share to increase gradually without requiring a new deployment for each increment. Listing 3 shows an illustrative feature flag configuration for a migrated route.

```
// migration-flags.config.ts (illustrative excerpt)
interface MigrationFlag {
  routeId: string;
  rolloutPercentage: number;
  healthStatus: 'healthy' | 'degraded' | 'rolled_back';
}

function isMigrated(routeId: string): boolean {
  const flag = getMigrationFlag(routeId);
  if (flag.healthStatus === 'rolled_back') return false;
  return hashUserId(currentUserId()) % 100 < flag.rolloutPercentage;
}
```

Listing 3. An illustrative feature flag check determining whether a given user reaches the migrated route, consulting the route's current health status (Hodgson, 2017).

4.5 Automated Codemods and Schematics

Structural changes that apply consistently across many components, such as converting an NgModule declared component to a standalone declaration, are applied through an automated codemod rather than by hand, following the schematics based approach Angular's own update tooling uses, so the transformation is applied identically everywhere it is needed rather than accumulating the inconsistencies described in Section 3.3. Listing 4 shows an illustrative codemod script.

```
# scripts/migrate-to-standalone.sh (illustrative excerpt)
npx ng generate @angular/core:standalone \
  --mode=convert-to-standalone \
  --path=src/app/account-settings

npx ng generate @angular/core:standalone \
  --mode=prune-ng-modules \
  --path=src/app/account-settings

npx eslint --fix src/app/account-settings
```

Listing 4. An illustrative automated codemod converting a module's components to standalone declarations consistently across every affected file (Angular, 2024; Gechev, 2022a).

4.6 Canary Release and Automated Rollback Gate

Before a migrated route's rollout percentage is increased, the pipeline checks a defined set of health signals, error rate, latency, and a small set of business specific success signals, against the traffic already routed to the migrated implementation, following the canary releasing pattern discussed in Section 2.3, and automatically reduces the rollout percentage back to zero if a signal crosses its defined threshold, rather than waiting for a human to notice and intervene manually. Listing 5 shows an illustrative continuous integration configuration enforcing this gate.

```
# .github/workflows/migration-canary-gate.yml (illustrative excerpt)
jobs:
  canary-check:
    runs-on: ubuntu-latest
```



steps:

- name: Query error rate and latency for migrated route
run: node scripts/query-canary-health.js --route=account-settings
- name: Evaluate health against thresholds
run: node scripts/evaluate-health.js --max-error-rate=0.01
- name: Roll back automatically if unhealthy
run: node scripts/rollback-if-unhealthy.js --route=account-settings

Listing 5. An illustrative continuous integration configuration checking canary health and rolling back automatically if a threshold is crossed (Beyer et al., 2016).

4.7 Migration Telemetry and Dependency Tracking

Every route's migration status, rollout percentage, and canary health history is logged, alongside which of the application's dependencies have been upgraded to a version compatible with the target architecture, following the semantic versioning tracking discussed in Section 2.7, giving the team a single, current view of overall migration progress rather than needing to reconstruct that view manually from scattered deployment records.

4.8 Common Pitfalls and Anti Patterns

Teams adopting the pipeline in Section 4 tend to encounter a small, recurring set of pitfalls, each of which quietly reintroduces one of the problems in Section 3 if left unaddressed. This subsection lists the three most common ones observed while designing the illustrative case study in Section 5.

The first pitfall is choosing a migration unit in Section 4.2 that is too large, for example an entire feature area spanning many routes, on the reasoning that migrating fewer, larger units finishes the overall migration sooner. A unit that is too large reintroduces the concentrated risk problem in Section 3.1 at a smaller scale than a full application rewrite but larger than the pipeline intends, since a defect anywhere in that large unit affects every route within it simultaneously.

The second pitfall is leaving the hybrid bootstrap from Section 4.3 in place long after every route has actually been migrated, on the reasoning that removing it can wait until other priorities are addressed. An unused hybrid bootstrap still adds bundle size and runtime overhead to every request, and the longer it remains, the more it accumulates as exactly the kind of deferred cleanup debt the migration was meant to resolve rather than relocate.

The third pitfall is a canary gate, described in Section 4.6, whose health thresholds are set once at initial configuration and never revisited as the application's normal baseline error rate and latency shift over time, whether due to unrelated application changes or seasonal traffic patterns. A stale threshold can cause the gate to either roll back a genuinely healthy migration unnecessarily or fail to catch a genuinely unhealthy one, depending on which direction the baseline has drifted.

V. ILLUSTRATIVE CASE STUDY

To make the pipeline's expected effect concrete, this section walks through an illustrative case study for a single production Angular application migrating from an AngularJS and early Angular hybrid codebase over twelve migration weeks, with the pipeline's elements adopted cumulatively in the order presented in Section 4. All figures in this section report simulated indicators constructed for illustration, not measurements from a production application, and should be read as a worked example of the kind of effect the pipeline is designed to produce rather than an empirical result.

5.1 Migration Progress by Cumulative Adoption Stage

The first indicator tracks the share of the application's routes fully migrated to the next generation Angular architecture, at each cumulative adoption stage, starting from a baseline of planning activity alone before any part of the pipeline was adopted.




Figure 3. Illustrative share of application routes migrated to next generation Angular, by cumulative pipeline adoption stage.

In the illustrative scenario in Figure 3, the strangler fig cutover strategy produces the first meaningful progress by giving the team a concrete, repeatable unit of migration work, automated codemods produce a further increase by removing the manual, error prone application of structural changes as a bottleneck, and the feature flagged rollout and canary gate together produce the largest remaining gains by making each route's full cutover a fast, low risk step rather than a cautious, manually monitored process.

5.2 Illustrative Incident Trend During Cutover Windows

The second indicator tracks user facing incidents specifically occurring during a cutover window, meaning the period when a migrated route's traffic share is being increased, reported per migration week across the same twelve week period.




Figure 4. Illustrative user facing incidents during cutover windows, before versus after the zero downtime migration pipeline.

The illustrative before series in Figure 4 fluctuates in the low single digits with no clear downward trend, consistent with incidents arising from ad hoc, manually monitored cutovers with no automated rollback mechanism, while the after series drops sharply following adoption and settles near zero, consistent with the canary gate in Section 4.6 catching and automatically rolling back an unhealthy rollout before it affects a meaningful share of users.



Table 3. Illustrative case study summary by adoption stage

Stage	Pipeline element added (Section 4)	Routes migrated, percent (illustrative)	Note
Stage 1	None, baseline	4	Planning activity only
Stage 2	4.2 Strangler fig cutover adopted	28	Concrete, repeatable migration unit established
Stage 3	4.5 Automated codemods adopted	55	Manual, inconsistent structural changes removed as a bottleneck
Stage 4	4.4 Feature flagged rollout adopted	79	Rollout percentage increased without a new deployment per step
Stage 5	Full pipeline with canary gate	98	Full cutover becomes fast and low risk per route

VI. RESULTS AND DISCUSSION

6.1 What the Pipeline Does Not Eliminate

Even at full adoption in Table 3, the pipeline does not eliminate the underlying effort of rewriting legacy code to the target architecture, and the results in Figure 3 should not be read as evidence that migration effort itself has been reduced to near zero. What the pipeline changes is where that effort's risk shows up, concentrated in a single big bang cutover without the pipeline, distributed across many small, individually low risk cutovers with it, following the continuous delivery principle discussed in Section 2.2, not the total amount of engineering work the migration requires.

This distinction matters for how a team plans a migration's timeline, since the pipeline's benefit is primarily in reduced risk and reduced incident exposure during cutover, illustrated in Figure 4, rather than in a proportionally faster overall migration; a team expecting the pipeline itself to substantially shorten the calendar time a large migration requires is likely to be disappointed, since most of that calendar time is still consumed by the underlying code migration work in Sections 4.3 through 4.5, not by the cutover mechanism in Section 4.6.

6.2 Cost of Pipeline Overhead at Small Scale

The hybrid bootstrap, feature flagging infrastructure, canary health monitoring, and telemetry described in Section 4 are not free to configure and operate, and a small application with low traffic volume or a short remaining lifespan is unlikely to recover that cost at the same rate the illustrative case study in Section 5 reports, since canary releasing in particular depends on enough traffic volume for a small percentage rollout to produce a statistically meaningful health signal, an assumption that does not hold well for a low traffic application.

A useful, if informal, signal for when the overhead in Table 3 is worth paying in full is the application's expected remaining operational lifespan weighed against its current user base size, following the general logic in Section 2.3 that canary releasing's value scales with the traffic available to canary against; a small, low traffic internal tool nearing planned retirement may reasonably accept a shorter maintenance window and a simpler, more manual migration rather than adopting the full pipeline, while a large, long lived, high traffic application is precisely the case this article's pipeline is designed for.

6.3 Distinguishing a Genuine Migration Blocker from a Routine Canary Fluctuation

A recurring source of wasted effort during the illustrative case study's later stages, once the canary gate in Section 4.6 had been operating for several weeks, was a team's tendency to treat an ordinary, transient fluctuation in the canary's health signal as evidence of a genuine migration defect requiring investigation, rather than as expected noise around a



threshold that Section 4.8 already warns can grow stale. Table 5 summarizes an illustrative rule of thumb for separating the two cases, distinguishing a single check that briefly crosses a threshold from a sustained pattern across several consecutive checks.

Table 5. Illustrative rule of thumb for distinguishing a canary fluctuation from a genuine migration blocker

Observed pattern	Likely interpretation	Recommended response
Single check briefly crosses threshold, next check normal	Routine fluctuation, often traffic composition	No action, continue monitoring per Section 4.6
Two consecutive checks crossed, then recovers	Borderline, possibly threshold drift	Hold current rollout percentage, do not advance further
Three or more consecutive checks crossed	Likely genuine defect in the migrated route	Automated rollback per Listing 5, begin investigation
Crossing correlates with a known traffic spike	Capacity related, not a code defect	Review infrastructure scaling before resuming rollout

This distinction matters because an overly sensitive rollback policy, one that treats every single crossed check as a genuine blocker, erodes a team's trust in the canary gate over time in much the same way a chronically noisy alerting system erodes trust in monitoring generally, following the alert fatigue concern noted in Section 2.3, eventually leading teams to raise thresholds so high that the gate stops catching real problems at all.

VII. LIMITATIONS

The case study in Section 5 uses simulated indicators constructed for illustration and does not report measurements from a production application, so the specific migration progress and incident figures shown in Figures 3 and 4 should be read as illustrative of the kind of effect the pipeline is designed to produce rather than as an empirical claim. The pipeline also assumes an application architecture, route based code splitting and a hybrid bootstrap mechanism, that fits a typical Angular application's structure but may not translate directly to an application with an unusually monolithic routing structure, which would need to define its own bounded migration units before Section 4.2's strategy could apply as described. Finally, the codemod example in Listing 4 assumes Angular's own standalone conversion schematic covers the specific structural change in question, which will not be true for every structural change a given application's migration requires, some of which will need a custom, hand written codemod rather than an off the shelf one.

VIII. FUTURE WORK

Three directions follow from the limitations in Section 7. First, an empirical study measuring the pipeline's effect on a production Angular application, rather than through simulated indicators, would test whether the illustrative migration progress and incident trends in Section 5 hold in practice and at what adoption cost. Second, a more systematic treatment of which structural migration changes are amenable to an off the shelf schematic versus requiring a custom codemod, extending the limitation noted in Section 7, would help teams estimate the codemod development effort a given migration actually requires before committing to a timeline.

Third, the canary health threshold discussed in Section 4.6 was presented as a fixed, manually configured value; a follow up study on adaptive thresholds that adjust automatically as an application's baseline error rate and latency shift over time would address the staleness pitfall described in Section 4.8 more robustly than a periodic manual review alone.

Fourth, the rule of thumb in Table 5 for distinguishing a routine canary fluctuation from a genuine migration blocker was constructed informally from the illustrative case study's later stages rather than derived from a statistically grounded



model; future work could formalize this distinction as a proper anomaly detection method applied to the canary health signal, potentially reducing both the false alarm rate and the risk of missing a genuine defect that a simple consecutive check count might overlook.

IX. CONCLUSION

This article presented a zero downtime migration pipeline for legacy Angular applications, built around four coordinated elements, an incremental, route level cutover following the strangler fig pattern, dual version coexistence through a hybrid bootstrap, automated codemods and schematics for structural changes, and a feature flagged, canary gated rollout with automated rollback. An illustrative case study for one production Angular application suggested the combined elements can substantially increase migration progress while reducing cutover related incidents, though these results are illustrative rather than empirical. The pipeline's central insight, that a migration's total effort is largely fixed by the amount of legacy code that needs rewriting, but the risk that effort carries during cutover is not fixed and can be substantially reduced through incremental, automatically verified rollout, reflects a broader point this article's discussion in Section 6.1 returns to, that zero downtime migration is fundamentally a risk management discipline applied to an already necessary body of engineering work, not a technique for avoiding that work.

Appendix A: Zero Downtime Migration Adoption Checklist

This checklist summarizes the practical steps a team can use when adopting the pipeline described in Section 4, organized by the same four elements from Section 4.1. It is offered as a starting point rather than an exhaustive audit.

A.1 Strangler fig cutover (Section 4.2)

1. Identify a bounded migration unit, following the routing example in Listing 1, sized to keep risk contained per the correction to the anti pattern in Section 4.8.
2. Confirm each migration unit has a clear success definition before starting its migration, not only a vague sense that it has been converted.
3. Sequence migration units by risk and traffic volume, migrating a lower risk unit first to validate the pipeline before attempting a higher risk one.
4. Track the cumulative migration progress indicator, following Table 3, to confirm units are actually completing rather than remaining perpetually in progress.

A.2 Hybrid bootstrap coexistence (Section 4.3)

1. Confirm the hybrid bootstrap configuration in Listing 2 correctly upgrades and downgrades components at the boundary between legacy and migrated code.
2. Document which specific routes or modules remain on the legacy framework at any point in time, avoiding ambiguity about the coexistence boundary.
3. Remove the hybrid bootstrap promptly once the last route is migrated, following the correction to the anti pattern in Section 4.8.
4. Monitor bundle size and runtime overhead attributable to the hybrid bootstrap itself, to have a concrete number motivating its timely removal.

A.3 Automated codemods (Section 4.5)

1. Confirm an off the shelf schematic, following Listing 4, actually covers the specific structural change in question before assuming it will work unmodified.
2. Run the codemod against a representative sample of the codebase first, reviewing the output before applying it broadly.
3. Pair every codemod run with the project's existing linting and formatting tools, following Listing 4's example, to catch a codemod introduced inconsistency immediately.
4. Document any custom codemod written for this migration for reuse in a similar future migration.

A.4 Feature flagged canary rollout (Sections 4.4, 4.6)

1. Confirm the feature flag check in Listing 3 consults the route's current health status, not only its configured rollout percentage.
2. Confirm the canary health thresholds in Listing 5 reflect the application's actual current baseline, revisiting them periodically per the correction in Section 4.8.
3. Confirm an automated rollback actually reduces the rollout percentage to zero, not merely alerts a human to consider doing so.



4. Track the incident indicator in Table 3 by migration week to confirm the canary gate is holding, not only at initial adoption.

Appendix B: Illustrative Route Migration Lifecycle Example

This appendix works through a single illustrative route migration end to end, from selection through the pipeline's four elements to full cutover, to make the abstract description in Section 4 concrete.

Table 6. Illustrative route migration, from selection to full cutover

Stage	Illustrative content
Migration unit selected (Section 4.2)	The account settings route, moderate traffic, low business criticality
Hybrid bootstrap in place (Section 4.3)	AngularJS and Angular coexist, this route not yet upgraded
Codemod applied (Section 4.5)	Component converted to standalone declaration, following Listing 4
Canary rollout begins (Section 4.4)	Feature flag set to 5 percent of traffic, health monitored
Health check (Section 4.6)	Error rate and latency remain within threshold across three checks
Outcome	Rollout percentage increased to 100 percent, route fully migrated, telemetry updated per Section 4.7

Listing 6 shows the illustrative telemetry entry this route migration would produce, following the fields described in Section 4.7, and Table 6 summarizes the stage by stage progression this single route passed through.

```
// Illustrative migration telemetry entry, following Section 4.7
{
  routeId: 'account-settings',
  migrationStatus: 'fully_migrated',
  rolloutHistory: [5, 25, 50, 100],
  canaryHealthChecks: 3,
  rollbacksTriggered: 0,
  fullCutoverDate: '2026-02-14',
}
```

Listing 6. An illustrative migration telemetry entry for the route worked through in Table 6 (Section 4.7).

Appendix C: Illustrative Migration Responsibility Matrix

This appendix summarizes an illustrative division of responsibility across the roles typically involved in operating the pipeline described in Section 4, to make concrete who is expected to act at each stage of a route's migration. Table 7 uses a simple designation of primary responsibility, secondary involvement, or no direct involvement for four illustrative roles.



Table 7. Illustrative migration responsibility matrix by pipeline stage and role

Pipeline stage (Section 4)	Feature team engineer	Platform or infrastructure engineer	Engineering manager	On call responder
4.2 Strangler fig cutover selection	Primary	Secondary	Secondary	None
4.3 Hybrid bootstrap maintenance	Secondary	Primary	None	Secondary
4.4 Feature flagged rollout control	Primary	Secondary	Secondary	None
4.5 Automated codemod authorship	Primary	Secondary	None	None
4.6 Canary gate configuration	Secondary	Primary	None	Secondary
4.6 Automated rollback response	Secondary	Secondary	None	Primary
4.7 Migration telemetry review	Secondary	Primary	Primary	None

A team adopting the pipeline at a smaller scale, where the platform or infrastructure engineer and on call responder roles in Table 7 may be filled by the same one or two people rather than a dedicated separate team, can still use the same division of responsibility as a checklist for what needs covering, even if fewer distinct individuals cover it, following the small scale overhead discussion in Section 6.2.

Statements and Declarations

Funding. This work received no specific grant from any funding agency in the public, commercial, or not for profit sectors. [Funding statement placeholder, to be completed with actual funding information prior to submission.]

Conflicts of interest. The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article. [Conflicts of interest statement placeholder.]

Data availability. All numerical results presented in Section 5 and Section 6 are simulated for illustration and are not derived from a production dataset. [Data availability statement placeholder, to be completed if empirical data is added in a future revision.]

Author contributions. [Author contributions placeholder, to be completed prior to submission, for example following the CRediT taxonomy for conceptualization, methodology, and writing.]

Acknowledgments. [Acknowledgments placeholder for colleagues, reviewers, or institutions who supported this work but do not meet authorship criteria.]

REFERENCES

- Angular. (2022). Discontinued Long Term Support for AngularJS. Angular Blog. <https://blog.angular.dev/discontinued-long-term-support-for-angularjs-cc066b82e65a>
- Angular. (2023). Upgrading from AngularJS to Angular. Angular documentation. <https://v17.angular.io/guide/upgrade>
- Angular. (2024). Update Guide. Angular documentation. <https://angular.dev/update-guide>
- Beyer, B., Jones, C., Petoff, J., and Murphy, N. R. (Eds.). (2016). Site Reliability Engineering: How Google Runs Production Systems. O'Reilly Media.
- Fowler, M. (2004). StranglerFigApplication. [martinfowler.com. https://martinfowler.com/bliki/StranglerFigApplication.html](https://martinfowler.com/bliki/StranglerFigApplication.html)
- Gechev, M. (2022a). Angular v15 Is Now Available. Angular Blog. <https://blog.angular.dev/angular-v15-is-now-available-df7be7f2f4c8>
- Hodgson, P. (2017). Feature Toggles (aka Feature Flags). [martinfowler.com. https://martinfowler.com/articles/feature-toggles.html](https://martinfowler.com/articles/feature-toggles.html)



8. Humble, J., and Farley, D. (2010). Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation. Addison Wesley.
9. Preston-Werner, T. (2013). Semantic Versioning 2.0.0. <https://semver.org/>
10. Raible, M. (2020). What Is Angular Ivy and Why Is It Awesome? Okta Developer Blog. <https://developer.okta.com/blog/2020/02/12/angular-ivy>