

From Development to Deployment: A Systematic Approach to Operationalizing Infrastructure Applications

Nikhita Kataria

nikhitakataria@gmail.com

Abstract:

Often engineers want to work on frontend product development without paying too much attention to the backend infrastructure. For some, it comes for free and is taken for granted. For instance, how many times have we really thought about the various 9's promised by Amazon or Azure compute that we spin up for personal or professional projects. In this paper we talk about development to deployment considerations for an infrastructure application which often must solve multiple challenges for user facing application and must be a one solution fits all in certain cases.

Keywords: Infrastructure, Operational Excellence, Telemetry, Deployments, Linting.

I. INTRODUCTION

A lot of engineers naturally gravitate toward building frontend features—the parts of a product users can see and touch. Infrastructure, on the other hand, often fades into the background. It's assumed to “just work.” Most of us don't stop to think about what it takes to keep that promise—like the “four nines” of uptime from AWS or Azure—when we launch a VM for a weekend project or deploy something critical at work. But building infrastructure software isn't simple. It needs to be reliable, flexible, and able to support all kinds of applications—sometimes with just one implementation. It's expected to be rock solid even when it's solving messy, complex problems under the hood. In this paper, we walk through what it really takes to get an infrastructure app from development into production: the decisions, the trade-offs, and the kind of thinking that ensures it holds up in the real world.

II. DEVELOPMENT CONSIDERATIONS

Infrastructure applications often take longer to reach operational maturity. That's usually because they're complex, need to work across many different environments, and don't always have clear or consistent processes in place to support them. While building infrastructure software, it needs to work reliably in surge traffic and at off business hours to make life easier down the road for code creators and debuggers.

A. *Designing for Real-World Operations*

During development here are a few considerations that are absolute must while building infrastructure applications:

1) *Deterministic Behavior*: If it works like ‘A’ with a certain input parameter, it should always work like ‘A’ as long as the input parameters remain same i.e be reliable.

2) *Idempotent*: If you are working on creating deployment infrastructure just a one-click retry should ensure the same operation happens and same state machine is exercised. If a deployment button is clicked 2 times, it should result in same deployment output.

3) *State Aware*: Building on idempotency the code should be state aware to be able to resume operations. Note the state awareness should not be built in the compute, it should be a stateless application and state should be maintained in underlying storage.

B. Writing Code That Can Grow and Survive

Often infrastructure code is written once and rarely refactored unless it's a complete rewrite of the functionality. It is essential to write code that is clean yet solves for multiple complexities and hence should be flexible. OOPs concepts play a big role while working on infrastructure applications. Key development considerations:

1) *Modular Design*: Let's assume you are working on writing a backend code that allows users to canary and then deploy code. The key business logic for this remains same as both entail deployment of 'a' code onto 'a' host. The code should be written as an interface and be modular.

2) *Separation of concern*: A big mistake we often make is not to know when to write configurations vs MACROS. Keeping configs separate from functional logic is essential. This way they would be easier to test, update and hand off to others.

3) *Testing for behavior and breaking your code*: Users should not be taken for granted while writing infrastructure code and every weird scenario should be tested because more often than not, what you think is weird will be exercised. While writing test cases, consider "what ifs" along with what should the right behavior be.

III. ENVIRONMENT STANDARDIZATION

One of the biggest headaches in infrastructure work is making sure your environments match—whether you're working in development, staging, or production. Without this, you'll hear the classic "It works on my machine" excuse way too often. Creating environments that are easy to reproduce and transparent takes the guesswork out, cuts down on mistakes, and makes teamwork way smoother.

A. Reproducibility with Containers

Containers like Docker have become the easiest way to package everything your app needs so it runs the same everywhere.

1) *No more surprises*: Containers bundle your code, dependencies, and even the operating system details into one package. So, whether it's on your laptop, a test server, or production, it's the same. No more random bugs from missing software or different versions.

2) *Keep things separate and peaceful*: Each container can have its own tools and dependencies, so you can run Python 3.12 in one and Python 3.8 in another without any conflicts.

3) *Test it all locally*: Want to spin up your entire production setup—databases, services, configs—right on your machine or in staging? Containers make that easy. It helps you debug faster and avoid surprises when you push to production.

B. Configuration as Code

It's not just about the code, as an infrastructure control plane author, it is essential to manage all levels of settings carefully. Treating configs like code brings clarity and keeps everyone on the same page.

1) *Track every change*: Instead of messing with settings manually or scribbling them down in docs, put your config files in Git. That way, you can see who changed what and when, and roll back if needed. Wondering why that timeout got bumped? Git blame will tell you.

2) *Everyone stays in the loop*: When configs live alongside code, the whole team can see what's running where. No more secret tweaks or "tribal knowledge." It also makes onboarding new people easier and troubleshooting faster because you can quickly spot any config issues. For instance we ran an experiment to detect drift in configuration by introducing small mismatches between canary and production environment and measured average detection time to be around 5 hours with 10% mismatch and with alerting less than 10 minutes. This concluded that configuration as code and observability can reduce mean time to detect significantly.

IV. DEPLOYMENT STRATEGIES

At times deploying changes to infrastructure comes with cyclic loops and can be daunting. These dependencies need to be categorized into ring architecture also referred to as tiers. Having a solid deployment

strategy helps you avoid those scary surprises, keeps your systems stable, and makes life easier for everyone on the team.

A. CI/CD Pipelines for Infrastructure Applications

Continuous Integration and continuous deployment is a standard practice across industry with a heavy focus on standardized automation. There are key outcomes that are missed as part of these pipelines. An exhaustive CI/CD pipeline must have following dimensions:

1) *Linting & Unit Tests*: The pipeline starts by giving your code a quick health check. Linters scan through it like a grammar checker, flagging any messy formatting or common slip-ups to keep things clean and consistent. Then come the unit tests; they zoom in on individual parts of your code to make sure each one works the way it should. Catching issues at this stage saves a ton of time later, so you're not stuck chasing bugs down the line. Table 1 compares and contrasts famous linters on a simple HTTP application written in python.

Linters	Key Feature	Linting errors captured	Strengths
Pylint	Comprehensive linting	Docstring, naming, structure warnings	Code quality and rule enforcement
Flake8	Style and light linting	Line length, basic style	Fast, widely adopted
Ruff	All-in-one fast linter	Combines Flake8, isort, etc.	Speed and wide rule coverage
MyPy	Type checking	Needs type hints	Type safety in statically typed code
isort	Import sorting	No issues unless disordered	Clean import organization
Bandit	Security analysis	Warns on weak HTTP patterns	Catch security bugs early
Black	Code formatting	Reformats code to standard style	Style consistency

Table 1: Comparison on industry standard linters

2) *Build & Package*: If the code passes the first tests, it's time to build it into a deployable package. Think of it like boxing up your app with everything it needs—code, libraries, and all—so it can run exactly the same anywhere, whether that's your laptop or a production server.

3) *Integration Tests*: Integrtrion checks validate that different parts of your system talk to each other properly i.e. with the new code all upstream and downstream services still produce reliable and expected results. Even if the individual pieces work fine, sometimes things break when they come together. These tests help catch those issues. With an experiment we detected that even with 100% unit test coverage but no integration tests, 3 out of 10 critical bugs were caught. Adding integration tests caught 9 out of 10 concluding integration tests in real environments are essential for production confidence.

4) *Promote to Staging*: Once everything looks good, the code moves to staging—a safe space that mimics production. This is where you run final checks, simulate real-world use, and make sure nothing's going to surprise you once it goes live.

5) *Manual or Automated Rollout*: Finally, the update gets rolled out to production. Depending on your team's comfort level, this can happen automatically if all tests pass, or after a manual approval step for extra peace of mind.

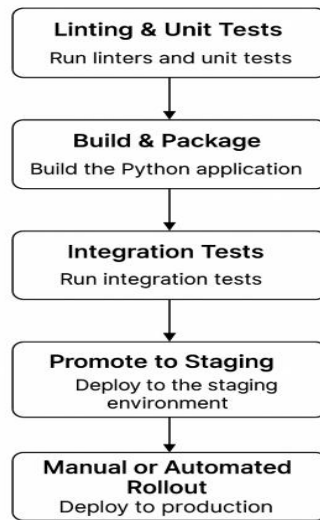


Fig. 1. Summarized outline of CI/CD pipeline

B. Safe Rollouts with Canary Deployments

While working on infrastructure or control plane it is important to ensure zero downtime which means the change management must be very well coordinated avoiding large set of changes being batched together. Canary deployments play a key role here to allow optimal tests to happen on a controlled set of hosts (or clusters or any logical grouping) which is a smaller percentage of entire fleet with following key benefits:

1) *Reduce the blast radius*: Rolling out a change to just a small group of users or servers is a smart way to catch issues early without causing a big mess. Think of it like trying out a new piece of software on a controlled set of employees gradually expanding to entire company and finally to the end customers and with every rollout iterating on the feedback allowing controlled fixes as well.

2) *Spot issues early*: By closely monitoring this canary group, you can detect issues—like bugs or performance drops—before they impact your entire user base. If something goes off-track, you can quickly pause or roll back the changes, reducing downtime and keeping the overall experience stable and reliable.

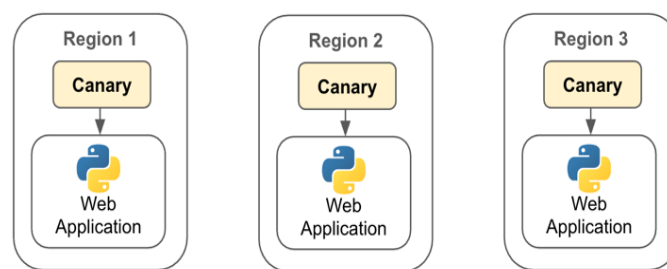


Fig. 2. Sample Rollout Plan

We compared two deployment strategies: all at once deployments vs a basic 10%-step canary rollout with a clear outcome that canary rollouts reduced both fault impact and recovery time as illustrated in Table 2.

Metric	All-at-Once	Canary (10%)
Error Rate Spike	6.3%	1.2%
Time to Detect Fault	240s	48s
Rollback Duration	120s	30s

Table 2. Metric comparison with and without canary.

V. OBSERVABILITY-FIRST DESIGN

Operationalization of infrastructure applications is expected to begin with visibility as infrastructure solves for multiple user end cases and can easily run into edge conditions which as often easily interpretable if applications are built with right observability in mind and not retrofitted based on run time outages. In this paper we strongly advocate for observability becoming a first-class priority 0 concern and integrated from early stages of development.

A. Metrics, Logs and Traces

It is essential to instrument infrastructure services with 3 core telemetry principles:

1) *Metrics*: Providing quantitative data on system performance and health with key metrics of latency, error rate, throughput queue sizes and retries visible into dashboards which are constantly monitored and reasoned on. Exploring solutions such as Prometheus based metrics which spans multiple dimensions with various levels of aggregations is advocated instead of having single dimension alerting to enable granular and high level views via the same metric.

2) *Logs*: Infrastructure applications most often have workflows built around actual data center hosts so having common schemas across services indexed on hosts, request ID's and service context is essential. Table 3. presents a comparison of a few standard centralized logging solutions tailored to a basic goLang application managing infrastructure around hosts with 50 concurrent threads sending 1000 log entries of 1MB each. Table 4. presents conclusions from experiments and their suitability per use case.

3) *Traces*: Traces are a key to visualize the entire lifecycle of requests from development to post deployment providing a solid entry for debugging. We examined porting our metrics to open telemetry for latency and throughput to generate distributed traces and saw a significant reduction in mean time to resolve for different issues.

Platform	Performance with (50 threads)	Operational Fit for Infra	Observations
OpenSearch + Fluent Bit	15MB/s	Built for structured ops logs	Best balance of scale, cost, and query features
ELK Stack (Logstash)	~8MB/s	Typically used in large infra platforms	Mature ecosystem, wide adoption
Fluentd + Elasticsearch	~10MB/s	Pluggable and flexible	Lighter than Logstash, customizable
Loki + Promtail	~18MB/s	Best for ephemeral/logging tailing	Very light, integrates with Grafana

Table 3. Comparison of platforms for centralized logging tailored to infrastructure application.

Use Case	Best Fit
Modern cloud-native infra management	OpenSearch + Fluent Bit
Legacy but rich observability stack	ELK Stack
Resource-constrained host agents	Loki + Promtail (log tailing only)

Table 4. Conclusion from experiments conducted

Issue Type	Traditional Logging	MTTR After OpenTelemetry (min)	Improvement (%)
Latency	90	30	66.7%
Error Rate	60	25	58.3%
Throughput Drop	90	30	66.7%
Resource Saturation	90	30	66.7%
DB Query Failures	90	20	77.8%

Table 5. Mean time to resolve for issues pre and post open telemetry.

B. Shift-Left on Telemetry

Infrastructure applications cannot afford to defer instrumentation to post development as most of the times a testing environment for other application is a production environment for control plane infrastructure applications. We advocate addition of telemetry during development and making it a common code review and CI gate practice. Constructing of key objectives alongside meaningful business metrics and embedding them into application logic goes a long way driving impact. Take a DNS provisioning service, for example. It doesn't just track basic HTTP metrics—it also reports things like how many records were successfully created, why failures happened (such as bad input or provider timeouts), and how TTL values are spread out. These specific, meaningful insights help teams diagnose issues much faster when the service is running in production.

C. Observability-Driven Rollouts

We make rollout decisions by listening closely to what our systems are telling us. During phased deployments (like canary releases) we keep a constant watch on key metrics to make sure everything's running smoothly. If error rates spike beyond a safe limit, we automatically roll things back. If latency starts to creep up, alerts go out and we hit pause. We also compute health scores based on how well things are performing overall, and those scores act as manual checkpoints where someone on the team decides whether it's safe to move forward. In contrast, traditional deployments often depend on synthetic tests or, worse, waiting for user complaints to surface issues. That kind of reactive approach doesn't cut it for infrastructure applications, where a small misstep can lead to widespread outages or misconfigurations. Instead, we take an observability-first approach: telemetry isn't just something we check after core features are developed, it is a continuous input actively guiding rollout decision. We build observability into every stage, using live data, alert thresholds, and known behavioral patterns to decide in real time whether to keep going, hold off, or roll back a deployment. Observability-driven rollouts flip the model: telemetry becomes the control plane for deployment decisions.

1) *Canary Release with Metric Validation*: In a canary rollout, we start by sending a small size of traffic to new version of a service. This gives us a safe window to watch how the new code behaves under real conditions. Throughout this phase, we closely monitor telemetry from the canary instance and compare it to what "normal" looks like. Take our DNS provisioning service as an example: we track high-percentile latency (like p95 and p99), error rates, and overall request volume. If we see latency spike above 500 ms, error rates climb past 1%, or traffic patterns dip unexpectedly, we treat that as a red flag. The system then automatically pauses or rolls back the deployment to avoid wider impact.

2) *Health Scoring as a Deployment Gate*: Traditional approach is to watch a dashboard of multiple metrics metric during deployments however a smarter approach is to combine metrics to generate an overall health score that blends latency, throughput, availability and error rates into one scope to help catch issues that might not otherwise stand out individually but become risky when combined.

3) *Observability Built into CI/CD*: Before rolling a change out to more users, we must run checks that query live metrics to see if things look stable. After a rollout, we don't just declare success but we look at trace samples and metric trends to catch any silent regressions. Embedding real-time dashboards into pull request and deployment view so that engineers can see what is happening and make informed decisions is a key aspect tying development and deployment closely.

4) *Turning Rollout Issues into Learning Moments*: Whenever something goes wrong during a rollout—whether it's just a brief pause or a full rollback—we make sure to record it in a shared tracker. We note what happened, what caused it, how it was fixed, and include screenshots from dashboards to give full context. It's not just for postmortems—it helps create a feedback loop. Developers can look back at past issues, recognize patterns, and use those insights to write better tests and add smarter telemetry in future deployments.

D. Operational Dashboards

Infrastructure applications in each environment need to have customized thresholds and dashboards built on multiple dimensions such as the one powered via Grafana. These dashboards should be version controlled, managed and most importantly reviewed before every update to ensure they stay current, and the configured thresholds are tuned according to the expected service load. Table 5 presents an experiment conducted with telemetry as a deployment gate goal to demonstrate how real-time metrics improve deployment safety on a scenario where a new software version introduces a misconfiguration leading to 3X latency increase demonstrating metrics-based gating leading to early fault detection and proactive rollbacks.

Metric	Metrics-Gated	No Gate
Time to Detect	10 min	35 min
Time to Rollback	12 min	45 min
User Tickets Raised	0	17

Table 5. Key Metrics improved with and without metrics built into deployment workflow

VI. CONCLUSION

Infrastructure applications often operate behind the scenes, yet they play a critical role in ensuring that user-facing products remain reliable, scalable, and performant. This paper has highlighted the unique challenges of developing, deploying, and operationalizing infrastructure software—from designing deterministic, idempotent, and modular code to enforcing environment consistency with containers and configuration-as-code practices. We outline how deployment strategies such as canary releases when coupled with observability inbuilt can significantly reduce outage risks and reduce time to detect and recover. Adding telemetry early in the development cycle and monitoring real time metrics during day-to-day deployment operations empower engineers to make safe ramps and data driven rollout decisions which enhances operational excellence. By ensuring telemetry is treated as a design principle, infrastructure applications become more resilient and easier to manage at scale. This systematic approach from development through deployment and monitoring ensures infrastructure software can met complex use case needs with high reliability and availability in all types of environments.

REFERENCES:

- [1] J. E. Smith and R. Nair, *Virtual Machines: Versatile Platforms for Systems and Processes*. Morgan Kaufmann, 2005.
- [2] OpenTelemetry, "OpenTelemetry Documentation," [Online]. Available: <https://opentelemetry.io/docs/>. [Accessed: Jun. 26, 2025].

- [3] Cloud Native Computing Foundation, "OpenTelemetry Project Overview," [Online]. Available: <https://www.cncf.io/projects/opentelemetry/>. [Accessed: Jun. 26, 2025].
- [4] Lightstep, "How Distributed Tracing Reduces MTTR," [Online]. Available: <https://lightstep.com/blog/how-distributed-tracing-reduces-mttr/>. [Accessed: Jun. 26, 2025].
- [5] New Relic, "Mean Time to Resolve (MTTR) and Observability," [Online]. Available: <https://newrelic.com/blog/best-practices/mean-time-to-resolve-mttr>. [Accessed: Jun. 26, 2025].
- [6] Google Cloud Blog, "Why Observability Matters for MTTR," [Online]. Available: <https://cloud.google.com/blog/products/observability/why-observability-matters-for-mttr>. [Accessed: Jun. 26, 2025].
- [7] M. Sigelman et al., "Dapper, a Large-Scale Distributed Systems Tracing Infrastructure," Google Research, 2010. [Online]. Available: <https://research.google/pubs/pub36356/>. [Accessed: Jun. 26, 2025].
- [8] Grafana Labs, "How Observability Improves MTTR," [Online]. Available: <https://grafana.com/blog/2021/01/14/how-observability-improves-mttr/>. [Accessed: Jun. 26, 2025].