

How Big Tech

Companies
do **Multiple
Releases** in
a day

2025

Executive Summary

Times are as critical as ever, companies are under relentless pressure to ship features faster and stay ahead of the competition. However, while engineering efforts race forward, critical processes like documentation, testing, and validation often lag behind — leading to broken release cycles, higher support costs, and a poor customer experience.

qAPI, a codeless API testing platform, was built to solve this very challenge. After facing growing pains firsthand during its global expansion in early 2023, qAPI recognized that traditional API testing

methods — dependent on slow scripts, fragile automations, and delayed validation — simply couldn't keep pace with modern development speeds.

As applications today promise near-zero downtime and 24/7 availability, the need for reliable, scalable API testing has never been greater. Companies that succeed are those who not only deploy faster but also validate faster — releasing stable features every few days, compared to the industry average of three weeks for feature completion and rollout.

This ebook outlines the realities behind today's API release struggles and reveals a better path forward: faster, smarter, and more resilient API deployment strategies, powered by the right tools and processes. Learn how leading teams are accelerating delivery, reducing defects, and future-proofing their digital services with practices built for the speed of modern innovation.



Page of Contents

Why did the Tech Giants shift towards an API-First Testing approach? Page 05

The key incident types Page 11

API incident attacks Page 15

Meta's API Testing and Release Process Page 19

Google's API testing and release process Page 22

Amazon API testing and release process Page 25

Netflix API control and deployments Page 28

Stripe API testing and release process Page 32

Why API Testing Needs a Strategic Reboot Page 36

API-First in 2025: The Strategy Behind High-Velocity Engineering Teams



Introduction to API Testing and Its Global Impact

APIs, or Application Programming Interfaces, are becoming increasingly critical and important aspects of modern software. They enable seamless communication between systems. Over the past decade, their importance has skyrocketed, driven by cloud computing, microservices, and digital transformation needs.

In 2025, speed is not a luxury—it's a necessity. Yet over **67% of enterprises struggle to release API updates more than once a week**, often due to brittle UI-based tests, siloed QA processes, and an over-reliance on manual testing. As digital experiences become increasingly API-driven, delays in testing and deploying APIs are directly impacting time-to-market, user satisfaction, and revenue growth.

According to the **World Quality Report 2024**, **58% of QA leaders state that they are unable to test APIs early in the lifecycle as a key bottleneck in achieving continuous delivery**. Adding to that the growing complexity of microservices, cloud-native architectures, and global compliance pressures—the need for a scalable,

automated shift-left testing strategy becomes undeniable.

APIs will likely become even more central in the next five years, with trends like AI-driven testing, enhanced security, and standardization shaping their evolution.

So how do industry leaders like Netflix, Amazon, and Facebook deploy changes hundreds of times daily without compromising on quality or breaking downstream systems?

The answer lies in their **API-first testing strategy**—an approach where testing starts early, focuses on business-critical risks, and is built for scale, speed, and collaboration.

They have mastered the art of rapid, reliable releases by adopting API-first testing, which ensures quality without slowing down innovation.

But achieving this level of agility isn't easy. As APIs grow in complexity and scale, so do the challenges: ensuring reliability, performance, and security while maintaining the speed. Traditional testing methods often fall short, leading to bottlenecks that hinder progress.

This is where API-first testing comes into play—a complete shift designed to push teams to test early, automate everything, and integrate seamlessly into CI/CD pipelines.

In this guide, we'll dive deep into how big tech companies have harnessed the

power of multiple daily releases to stay competitive and deliver exceptional user experiences. You'll learn:

- How industry giants like Stripe, Netflix, and Amazon use API-first testing to accelerate deployments.
- The tools, strategies, and best practices they rely on to ensure quality at speed.
- The step-by-step process to implement these approaches in your organization.

Whether you're aiming for tangible outcomes like new revenue streams or intangible wins like building customer loyalty, mastering API-first testing is key to staying ahead in an API-driven world.

Why did the Tech Giants shift towards an API-First Testing approach?






Consumer Demand and Seamless Omnichannel User Experience

The increase in API adoption is heavily influenced by consumers who want better user experiences. As digital platforms increase, users expect seamless interactions, which is pushing the need for regular updates and improvements to APIs.

For instance, the shift towards accommodating assistive technologies requires companies to invest resources to ensure their APIs can support the user needs effectively.

Whether it's in personalized recommendations, one-click payments, or real-time tracking, product expectations have shifted to **"always on, always evolving."** Frequent API releases allow engineering teams to quickly deliver value, respond to user feedback, and stay ahead of competitors.



2

Speed and Frequency of Releases

Now development lifecycles in teams often demand multiple deployments per day, especially in customer-facing platforms where responsiveness drives retention and revenue. Clearly, traditional UI-based testing approaches are too slow and fragile to keep up with this pace.

APIs, on the other hand, provide stable, **well-defined contracts** that can be tested independently of the UI. This allows for:

- Faster validation of features before the frontend is built
- Rapid iterations without waiting on full-stack readiness
- Continuous delivery without compromising test coverage

3

Early Bug Detection and Shift-Left Testing

Detecting bugs early in the development cycle is significantly cheaper and less disruptive than finding them in staging or production.

API-first testing enables this **“shift-left” approach** by allowing engineers to validate business logic and service behavior as soon as the API contract is defined. This ensures:

- Issues are caught before integration or UI testing stages
- Developers and testers can collaborate earlier and more effectively
- Teams avoid expensive rework caused by late-stage failures

The earlier the feedback loop, the stronger the product quality—and API-first testing is what makes that early loop possible.

4

Decoupling Frontend and Backend Workflows

For any software team, no one wants to work on monolithic applications. Services are modular, teams are cross-functional, so it makes sense that the **frontend and backend are developed in parallel**. With an API-first approach:

- Backend teams can build and test services independently
- Frontend teams can consume mock APIs or defined schemas to facilitate development until APIs are ready for integration
- Reduces dependency and coordination overhead between teams

This parallelization leads to faster release cycles and fewer integration issues—a much needed option in today’s distributed development models.

5




Scaling Testing with Microservices




As organizations adopt microservices, the number of deployable components can grow into the hundreds or thousands. Testing this scale through the UI is simply not feasible.

API-first testing solves this by:

- Making each service operational in isolation
- Creating reusable test suites across services
- Centralizing testing at the service layer where changes occur most frequently

It brings the needed scalability, modularity, and efficiency to QA processes in ways that UI-based tests just cannot.





6

Reducing Maintenance and Improving Test Stability

UI tests are notoriously fragile—minor visual changes can break entire test suites. This leads to increased maintenance, test flakiness, and wasted engineering hours.

In contrast, APIs are versioned, consistent, and contract-driven. Testing them directly:

- Reduces false positives and test flakiness
- Simplifies maintenance by focusing on logical behavior rather than visual structure
- Ensures tests remain valid across iterations of UI or frontend redesigns

The result is lower **QA overhead and greater test reliability** over time.

7

Faster Feedback Loops Through CI/CD Integration

To achieve continuous integration and continuous delivery, teams need **immediate feedback** on whether a code change has broken something critical.

API-first testing integrates cleanly with CI/CD pipelines, allowing:

- Tests to be triggered automatically on every commit or merge
- Developers to get results within minutes instead of hours
- Teams to catch regressions before they affect downstream systems

This leads to **greater deployment confidence**, fewer rollbacks, and higher overall velocity.



Built-In Security and Regulatory Compliance

APIs are the invisible secure layer of digital systems—and increasingly they are becoming the target of attacks and regulatory scrutiny. Security and compliance requirements are growing across all industries, especially in finance, healthcare, and e-commerce.

API-first testing helps embed security early by:

- Validating authentication, authorization, and input data at the API level
- Ensuring consistent behavior under compliance frameworks like PCI-DSS, GDPR, and HIPAA
- Making these validations part of the regular build and release cycle

It reduces the risk of data breaches, regulatory fines, and reputational damage by making **security testable and repeatable**.




Stronger Observability and Quality Insights

As systems become more distributed, it becomes critical to understand how they behave across multiple dimensions—functionality, performance, stability, and process flow.

API-first testing enables deeper observability by:

- Tracking coverage at the service and endpoint level
- Generating real-time metrics like response time, error rates, and pass/fail trends
- Enabling teams to correlate test results with production behavior through integrated dashboards

This kind of visibility empowers teams to make data-informed decisions about quality, reliability, and performance.



10

Internal Efficiency and Agility

UI tests are notoriously fragile—minor visual changes can break entire test suites. This leads to increased maintenance, test flakiness, and wasted engineering hours.

Speaking of quickly delivering value for some organizations, frequent API releases are critical for maintaining internal agility. As larger tech companies are often viewed as slower due to their size, having a rapid release cycle allows them to pivot quickly in response to changing market dynamics and internal requirements.


This adaptability is only built upon a foundation of maintainable and readable code, ensuring that while speed is prioritized, the integrity of the codebase is not affected.

Regular releases also create opportunities for developers to address any setbacks missed during expedited deployments, thereby creating a sustainable development cycle.

As our infrastructure grows more complex, high-profile outages, untested dependencies, and mounting compliance requirements have forced companies to rethink their release strategies. What once worked for you with slower monolithic software cycles is now causing problems in the time of high availability, real-time data, and distributed services.



The key incident types that catalyzed the global shift toward incremental, automated API releases.



High-Impact Outages from Monolithic Releases

The Problem: Centralized Risk, Disproportionate Impact

Monolithic architectures combine all services, components, and dependencies into a single unit. While it feels convenient in early stages of development, they can introduce serious risks as systems start to scale:

The modern API must serve a global, multi-device world—phones, TVs, wearables, third-party apps. Each expects data in real time, personalized, and always available. No exceptions.

- Mobile apps today make **200+ API calls per user session** (Postman, 2023).
- Over **83% of global internet traffic** is now API-driven (Akamai, 2023).
- Even a **100ms latency increase** can lead to a **1.2% drop in user retention** (Cloudflare, 2023).

When Twitter's API went down for 3 hours in 2023, it disrupted over 11,000 connected apps—from social logins to publishing tools. The ripple effect? A measurable 18% drop in user retention across key partners.

What's changing:


Incremental API deployments allow teams to patch issues, release improvements, or test features without bringing the entire system down—or impacting millions at once.

Why Incremental API Releases Matter

- **Reduced Blast Radius:** Releasing APIs as independently testable and deployable units prevents systemic collapse.
- **Faster Recovery:** Problematic endpoints or services can be rolled back without taking down the entire system.
- **Modern Practices:** Feature flags and microservices are now standard for isolating deployment risks.

2

Regression Bugs and Untested Interactions



The Problem: Latent Defects in Infrequent, Large Releases

When organizations defer API updates, they often accumulate tangled dependencies and risk breaking core functionality. The complexity of inter-service interactions grows, while test coverage often remains static or incomplete.

Key risks include:

- **Breaking Changes:** A single change can invalidate assumptions across services.
- **Rollback Headaches:** Monolithic rollbacks must undo multiple changes simultaneously—often introducing new bugs.
- **Delayed Feedback:** Failures are discovered late, costing time and user trust.

A single flaw could delay an entire launch or—worse—bring down live systems.

- Teams utilizing incremental API releases experience a significantly reduced mean time to repair (MTTR), **22 minutes on average**
- Take [Atlassian's Jira](#) outage in 2022: a bug in a shared authentication API took down customer access for four days, costing over \$15 million in refunds. Since then, they've migrated to scoped endpoints and decoupled deployments. Every API change now undergoes automated regression testing before shipping to production.


What's changing:

With microservice-based APIs and versioned contracts, companies can deploy independently, fix fast, and release daily—without waiting for the next “big release train.”

“Smaller releases allow iterative testing and faster recovery.”

Why Frequent Releases Are Safer

- **Real-Time Detection:** Smaller API releases are easier to test through continuous integration suites.
- **Agile Mitigation:** Incremental deployments can be paused, patched, or rolled back quickly.
- **Industry Benchmark:** Google executes over **4,000+ deployments daily**, maintaining **99.9% availability** thanks to modular API updates and automated regression testing.



Security & Compliance Pressures

APIs are now the #1 attack vector for enterprises. Between 2020 and 2023, API-related breaches surged 221% (Salt Security).

These aren't just theoretical vulnerabilities—they're expensive.

- The T-Mobile API breach (2023) leaked 37M customer records.
- The Optus API leak (2022) exposed 9.8M records, costing the company \$1.5M in regulatory fines.
- The average API breach now costs companies \$4.7M (IBM, 2023).


Why the spike? Traditional APIs were updated after long durations, with slow rollout of security patches. Today, best-in-class teams push updates weekly or daily patching vulnerabilities in hours, not quarters.

Delaying patches—even by days—can lead to:

- **Massive Fines:** Under GDPR, organizations face penalties up to €20 million or 4% of global revenue.
- **Data Breaches:** Unpatched APIs become attack vectors for malicious actors.
- **Audit Failures:** Monolithic systems lack the traceability regulators demand.

The Case for Incremental Security Releases

- **Rapid Patching:** Vulnerability fixes can be pushed to production within hours, not weeks.
- **Better Auditing:** Small releases create fine-grained audit logs for regulatory compliance.
- **Industry Trend:** According to **Gartner (2023)**, **73% of enterprises** now integrate API security directly into CI/CD pipelines—marking a clear shift from reactive to proactive security.



API incident attacks over the years





Tesla

- API flaw allowed unauthorized access to vehicle controls via third-party apps.
- Safety concerns; reputational damage.

2021



Zoom

- API vulnerabilities exposed meeting passwords and participant data.
- \$85 million settlement; reputational harm.

2021



British Airways

- API flaw led to a breach exposing 400,000 customers' payment details.
- £20 million GDPR fine; reputational damage.

2018



Experian

- API breach exposed credit scores and personal data of millions of South Africans.
- \$10 million settlement; reputational damage.

2020



Target

- API vulnerability led to a breach compromising 40 million credit card records.
- \$18.5 million settlement; reputational damage.

2013



Sony PlayStation Network

- API flaw led to a breach affecting 77 million users.
- \$15 million settlement; reputational damage.

2011



LinkedIn

- API flaw allowed scraping of 700 million users' data.
- Reputational damage; regulatory scrutiny.

2021



PayPal

- API flaw exposed sensitive user data during transactions.
- Reputational damage; potential GDPR fines.

2021



eBay

- API misconfiguration allowed unauthorized access to user account details.
- Customer trust loss; reputational harm.

2021



WhatsApp

- API vulnerability enabled attackers to scrape phone numbers of users.
- Privacy concerns; regulatory scrutiny under GDPR.

2021



DoorDash

- API flaw leaked delivery driver locations and order details.
- Reputational damage; customer complaints.

2021



Starbucks

- API outage disrupted mobile app payments and loyalty rewards.
- Lost revenue; customer dissatisfaction.

2023



Walmart

- API vulnerability exposed product pricing and inventory data.
- Competitive disadvantage; reputational harm.

2022



Spotify

- API bug caused playlist syncing issues across devices.
- User frustration; reputational damage.

2021



GitHub

- API rate-limiting issue disrupted CI/CD pipelines for developers.
- Developer productivity loss; reputational harm.

2021



Instagram

- API flaw allowed scraping of user profiles and private posts.
- Privacy violations; regulatory scrutiny.

2021



Adobe

- API vulnerability exposed Creative Cloud subscription data.
- Reputational damage; customer trust loss.

2021



Okta

- API breach exposed authentication tokens for enterprise customers.
- \$1 million in mitigation costs; reputational harm.

2022



Shopify

- API flaw leaked merchant data, including email addresses and order histories.
- Reputational damage; regulatory scrutiny.

2020



Expedia

- API outage disrupted hotel booking services globally.
- Lost bookings; reputational harm.

2022



Twilio

- API vulnerability exposed SMS message content and metadata.
- Privacy concerns; regulatory scrutiny.

2022



Salesforce

- API bug caused CRM data synchronization delays for enterprise clients.
- Operational inefficiencies; reputational harm.

2021



Netflix

- API flaw led to incorrect billing charges for some users.
- Customer complaints; reputational damage.

2021



Uber Eats

- API flaw exposed restaurant partner data, including menu pricing.
- Competitive disadvantage; reputational harm.

2022



Microsoft Azure

Coinbase

Fitbit

- API misconfiguration led to accidental exposure of customer storage accounts.
- Reputational damage; customer trust loss.

2023

- API vulnerability allowed unauthorized access to user wallet balances.
- Financial losses; reputational harm.

2022

- API flaw exposed user health data, including step counts and sleep patterns.
- Privacy concerns; regulatory scrutiny.





2021

Zero Downtime, Daily Deployments: What Makes It Possible for these Tech Giants

Releasing APIs multiple times per day might sound like a recipe for chaos—but for tech giants, it's the new normal. These companies aren't just fast; they're consistent, stable, and precise. So, what enables this level of release velocity?

It's not luck, and it's not just DevOps. It's a **disciplined API-first testing strategy**—designed to support automation, shift-left QA, and modular microservices at scale.

In this section, we break down the **core practices, tools, and cultural shifts** that enable leading tech firms to deploy APIs at lightning speed—without compromising quality, security, or user experience.



Meta's API Testing and Release Process



Code Commit & Branching Strategy

- **What Happens:** Engineers commit code either directly to the main branch for fast, continuous deployment or send it to feature branches that go through peer review and merging processes.

This dual-path system balances velocity and control—quick patches can be deployed instantly, while larger changes follow much longer and stricter review cycle.

The result: Developers can work in parallel without risking mainline stability, supporting Meta's need for high developer autonomy and velocity.



Three-Layer Continuous Integration (CI) Stack

- Next, every code commit activates a CI pipeline consisting of three progressively rigorous layers:
 - **Layer 1 – Build Validation:** Ensures that the code compiles successfully across all impacted services and products.
 - **Layer 2 – Static Analysis:** Uses linters and Meta's proprietary tool, Infer, to detect memory leaks, null pointer exceptions, and other critical issues without executing code.
 - **Layer 3 – Automated Testing:** Runs thousands of unit tests, integration tests, and end-to-end API validations.

The layered approach surfaces simple failures quickly while providing deep analysis as the code progresses through the stack.

Impact: Stops bad code early, reduces regressions, and increases confidence in deployment readiness.

3

Automated Intelligent Testing

Meta uses a **generic API testing tool**, a search-based testing tool, to generate and execute thousands of UI and functional test cases autonomously.

Manual test design is time-intensive and often misses edge cases. The tool uses machine learning to explore API and UI interactions that a human might overlook.

Meta Engineering reported that in just the first few months, the tool allowed engineers to **“detect and fix critical issues within hours—or even minutes—of code being written.”**

4

Conveyor: Using it Deploying Pipeline Automation

Conveyor is Meta’s internal tool that creates automated deployment pipelines for every service.

With 30,000+ deployment pipelines in active use (97% fully automated), Conveyor ensures every API, microservice, and backend component follows the same best-practice standards.

Impact: This drastically reduced the manual errors and standardized release quality across thousands of internal teams and systems.

5

Phased Rollout Strategy

How It Works: Once code passes CI, it enters a **three-phase rollout:**

- **Phase 1:** Deploy to Meta employees (dogfooding).
- **Phase 2:** Expand to 2% of production traffic.
- **Phase 3:** Full rollout to 100% of production.

This staged approach helps isolate failures early, ensuring issues don't impact the entire user base.

Impact: Engineers were able to monitor real usage in controlled environments and quickly halt rollouts if any unusual activity is detected



Health Check and Rollback Automation

After each rollout phase, a **Health Check Service** monitors:

- System performance (CPU, memory usage, crash reports)
- RPC performance (latency, error rates)
- Business metrics (engagement, conversion rates)

A rollback is triggered automatically if key metrics exceed error thresholds.

- **Impact:** This helped protect the user experience by avoiding a drop in ratings, even in high-speed deployment cycles.



Gatekeeper: Feature Flag Governance

Meta uses a **feature flag system called Gatekeeper** to toggle features on/off independently of deployment.

It allows decoupling code changes from user-facing feature releases. Engineers can safely deploy dormant code and activate it when ready.

- **Impact:** Supports A/B testing, progressive rollouts, and quick hotfixes without redeploying code.

Now let's look closely at how Google is able to handle daily multiple releases

Google: Building Safety into Velocity

Google deploys code thousands of times per day across its massive infrastructure — a scale that's only possible through tight integration of **automated API testing within its CI/CD pipelines**. Its approach emphasizes **early validation, mock simulation, and granular environment control**.

How does Google plan their deployment flow:



1

Local Simulation with Cloud Endpoints

Google begins API lifecycle workflows by simulating services **locally** using Cloud Endpoints Frameworks. Developers start by:

- Defining API specifications.
- Deploying locally with Endpoints and test the behavior of the API before pushing changes upstream.
- Catching errors in request/response handling, validation, and headers before even they are sent to staging.

2

API Proxy & Policy Enforcement via Apigee

Once an API passes local validation, it's promoted to the API management layer: **Apigee**. This isn't just an API gateway — it's a full traffic and policy enforcement system. Google uses it to:

- Apply authentication, quota limits, transformations, and backend routing rules.
- Run contract tests (e.g., is this version compatible with the last?).
- Provide a safe staging layer for testing how new changes behave in production-like conditions.
- Gradually shift traffic to the new version — using smart routing, traffic splitting, and real-time observability

3

Integrated Testing as Part of CI/C

Tests are integrated into **Cloud Build pipelines**, enabling:

- Triggered builds and deployments from Git repositories.
- Automatic regression tests after API changes.
- Rollbacks or version pinning based on test results — a critical part of safe multi-daily deployments.

Every commit can trigger:

- A full **build** → **test** → **deploy** pipeline.
- **Automated regression tests** across both unit and contract levels.
- **Rollback logic** based on test failures or behavioral deviations.
- **Version pinning** to isolate consumers of older APIs until they migrate to newer ones.

This infrastructure means that **APIs can change multiple times a day**, because they've made failure cheap to detect and easy to roll back.

Principle	What Google Does	What You Can Try
Shift left	Local API simulation with validation	Use tools like qAPI to simulate and validate
Manage exposure	Apigee as an enforcement layer	Introduce an API gateway with policy controls
Automate trust	Tests built into CI pipelines	Integrate contract and performance tests into CI/CD
Support rollbac	Auto reverts and versioning	Use Git-based deploys with rollback hooks or semantic versioning

Case Study: Amazon's API Testing Strategy

We tracked Amazon's plan for launching a new feature for its shopping service. To learn how they pushed updates to their product's API. They automated API tests using their collections within Code Pipeline, running functional, load, and performance tests with every code commit.


What happened next is it ensured the API can handle thousands of concurrent requests during peak times like Black Friday, with canary deployments testing new versions on a small user base before full rollout (A/B testing - AWS). We can clearly see this approach minimizes risks, enabling daily deployments while maintaining quality.

How did they manage that?

Small, autonomous teams (5-10 people) own the entire lifecycle of their services.


As stated in AWS Executive Insights, "Two-pizza teams do not hand over something they've launched to another team to run. This single-threaded ownership extends across the full customer experience."

It helped eliminate the handoffs and coordination overhead between teams, accelerating the development-to-deployment cycle.



1

API Design & Management via Amazon API Gateway



Developers begin by defining their REST or HTTP APIs in Amazon API Gateway. This includes:

- Method definitions and resource paths.
- Request/response models.
- Stage variables and throttling configs. These definitions can be tested via the Gateway console, simulating actual method calls and headers.

2

Automated Testing with Postman + CodeBuild

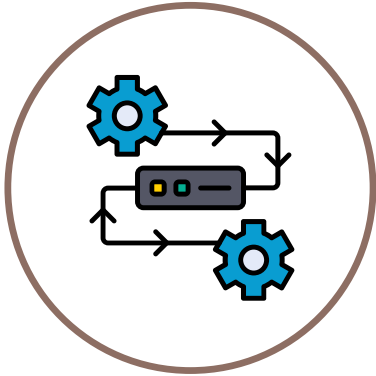


Amazon integrates their collections directly into **AWS CodeBuild**. This means:

- On each pull request or code commit, Postman tests are automatically triggered.
- These tests cover:
 - Contract validation
 - Expected error cases
 - Response times
 - Data integrity and schema consistency
- If even one test fails, the pipeline stops—the feature won't make it to staging.

3

Pipeline-Integrated QA with CodePipeline



Amazon then uses **CodePipeline** to plan deployments, with needed gates at each stage:

- After integration tests, APIs move to staging.
- Canary deployments are triggered with monitoring thresholds.
- Final production rollout is contingent on success metrics and alarms.

The Impact: What Changed After Adopting This Approach?

Metric	Before	After Adopting API Testing in CI/CD
Deployment Frequency	Weekly or bi-weekly	Every 11.6 seconds (avg.)
Rollback Rate	Moderate to high	Extremely low (due to canary testing)
Black Friday System Stability	Risk of downtime or lag	Handles millions of requests reliably
Time-to-Fix in Production	Hours to days	Minutes (due to observability + alerts)

Netflix – Scaling with Federated APIs and Controlled Deployments

Overview:

Netflix ships changes to its APIs and services thousands of times per day, supporting an enormous variety of client devices and global users. This velocity is possible through a combination of architectural shifts, federated API ownership, and highly automated release workflows.

The Deployment Strategy



1

Early Days: REST, But Fragmented

Initially, Netflix's APIs were RESTful and tightly coupled to client needs. Each device (TVs, phones, web) had its own custom-tailored endpoints, leading to:

- Duplicated logic across APIs
- Slower development due to inter-team coordination
- Difficulty scaling feature development across platforms

The Starting Point:

Netflix launched the "API.next" initiative to solve three key issues:

- Tight frontend-backend coupling
- API sprawl with many custom endpoints
- Inflexibility in evolving APIs without breaking clients

They needed an architecture that enabled modularity, faster iteration, and frontend/backend independence.

What they needed:

- A scalable, flexible API layer that allowed independent development.
- **Better coordination** between frontend and backend teams.
- A way to **evolve APIs safely** without breaking existing consumers.

Thus, they migrated to GraphQL

2

Why GraphQL?

Netflix adopted GraphQL to:

- Allow clients (like frontend apps) to **query only the data they need**.
- Reduce **over-fetching and under-fetching** that happened with REST.
- Offer a single endpoint rather than many REST routes.

Initial benefit:

- GraphQL acted as a unified layer between many backend services and the frontend — but the GraphQL schema and resolvers were still **centrally managed**, which became a bottleneck.

Enter Federation (via Apollo Federation)

To address the bottleneck of a centralized GraphQL server, Netflix adopted **Apollo Federation**, which allows:

a. Subgraphs (Microservices)

- Each backend team manages its own **subgraph** – a slice of the GraphQL schema specific to their domain (e.g., User, Video, Recommendations).
- These are essentially microservices that expose **GraphQL** schemas.

b. Supergraph

- All subgraphs are composed into one unified schema, known as the **supergraph**.
- The supergraph is what the frontend team queries against.
- Powered by the Apollo Router or Gateway, which routes incoming GraphQL queries to the right subgraph(s).

c. Independent Deployments

- Backend teams can **independently** update their part of the GraphQL schema (their subgraph).
- Since each subgraph is isolated and versioned, updates don't break the **supergraph** or the consumers (i.e., the frontend).

d. Safe Schema Evolution

- Apollo provides **schema validation tooling**, ensuring that changes are compatible with the supergraph and don't break existing queries.

Continuous Deployment + Testing

To make this work at scale:

- Netflix adds **automated schema checks** in CI/CD.
- **End-to-end testing** validates changes before they hit production.
- Teams deploy subgraphs often, enabling **rapid innovation** while minimizing risk.

3

Replay Testing & Canary Deployments

To ensure changes don't break production:

- Netflix uses **replay testing** to simulate real user traffic on new GraphQL APIs.
- Implements **sticky canary releases** to slowly expose a portion of traffic to the new codebase while preserving user session continuity.

4

Automated CD with Spinnaker


- Netflix uses [Spinnaker](#) for automated rollouts, supporting red/black and canary strategies.
- All deployments are baked as immutable AMIs (Amazon Machine Images), ensuring consistent behavior across regions and environments.
- Pipelines integrate with Jenkins and internal test suites before any code is allowed into production.

5

Workflow Orchestration with Conductor

- Complex backend processes are managed using Netflix Conductor, a microservice orchestration engine.
- This ensures business-critical flows (billing, personalization, etc.) continue running smoothly even as APIs evolve behind the scenes.

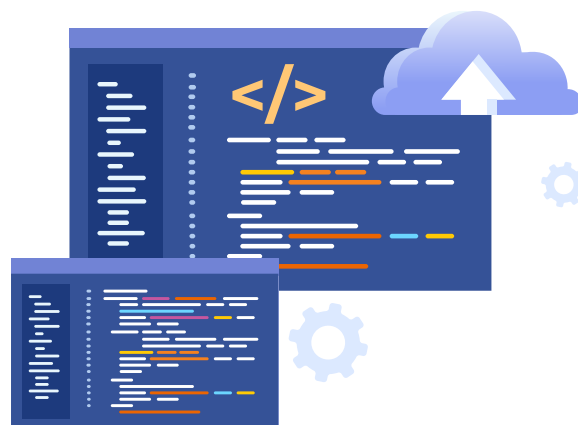
Takeaway: Netflix achieves high deployment velocity by *splitting API ownership, automating deployment pipelines, and testing in production with surgical precision.*



Case Study: Stripe – Precision APIs and Predictable Evolution

Overview:

For a global payments processor, reliability means trust. Stripe needed to ensure that APIs never break customer expectations, even with frequent updates. The shift was driven by:



- **Backward Compatibility:** Guarantee that even legacy integrations continue to function as APIs evolve.
- **Developer Trust:** Maintain a high standard for API quality, ensuring that clients don't face unexpected changes.
- **Risk Mitigation:** Minimize downtime and issues that could disrupt transactions across billions in payments.

The State Before the Shift:

- Early API deployments risked breaking changes that forced customers to continually re-integrate or face disrupted services.
- As API complexity increased, so did the risk of introducing regression bugs that could affect millions of transactions worldwide.

Chronology of Deployment Strategy



Immutable API Contracts with Date-Based Versioning

- Every Stripe API release is versioned by **date**, e.g., 2020-08-27.
- Consumers lock into specific versions, which **don't change** unless explicitly upgraded.
- Version override supported via the Stripe-Version header
- This ensures APIs evolve safely and predictably over time. ([Stripe Blog](#))



Progressive & Optional Upgrades

- New API features are added without breaking changes.
- Clients can choose when to adopt new features, making it easy for Stripe to ship changes daily without breaking or causing instability.



Automated Testing + Governance

Extensive Automated Testing: Each code change underwent hundreds of tests—from unit tests to contract tests that validate API responses against pre-defined models.

Governance Board: An internal API Governance Board review naming convention, error semantics, and payload standards before code is merged.

Linting and Schema Validation: Automated tools flagged any potential issues that could affect the public interface, ensuring a clean codebase before deployment.

4

Zero-Downtime Data Migrations

- Stripe runs massive data migrations using an in-house **Data Movement Platform**:
 - Supports chunked, online data movement across distributed systems
 - Guarantees no API downtime even when backend schemas are changing

Publicly available metrics indicate Stripe’s deployment frequency of core payment APIs was around 5,978 in 2022, averaging roughly 16.4 deployments per day—a pace that guarantees speed with stability.

5

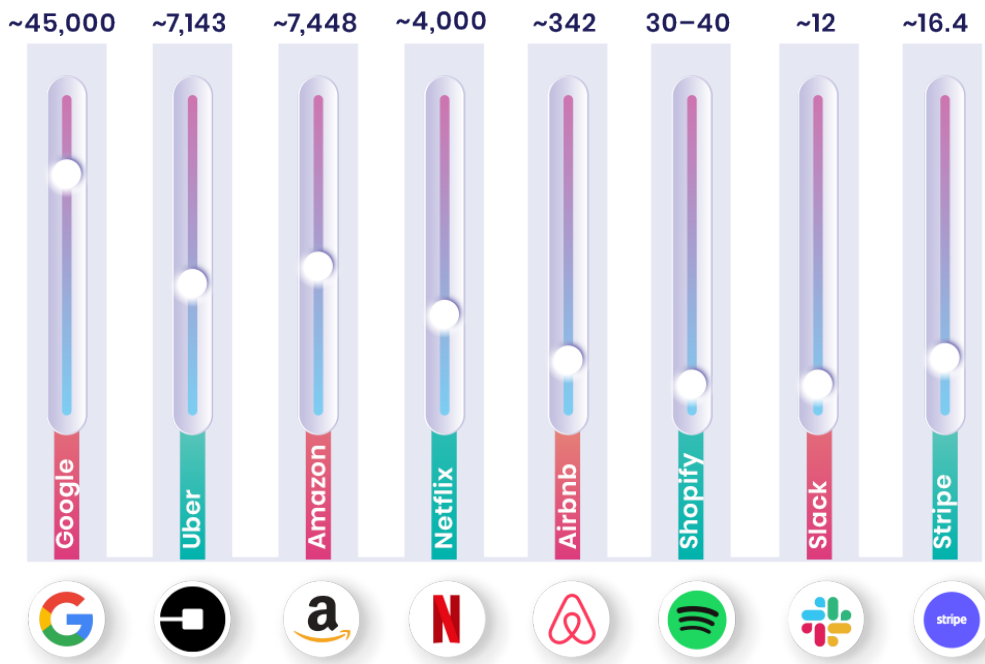
Semantic Releases & API Docs

- Releases are categorized as **minor, major, or breaking**.
- API changes automatically trigger doc updates, keeping developer documentation in sync with actual API behavior.
- Stripe’s ability to release daily comes from rigorous governance, **optional upgrades, and automation at every level of the API lifecycle**. <https://docs.stripe.com/upgrades>

API Release Strategy Comparison

Aspect	Google	Amazon	Netflix	Stripe
Release Frequency	Approximately 45,000 code commits per day, including 16,000 by developers and 24,000 by automated systems.	Averaging a deployment every 11.6 seconds, equating to roughly 7,448 deployments per day.	Over 4,000 deployments daily, facilitated by decentralized teams using in-house tools.	In 2022, executed 5,978 deployments of core payment APIs, averaging 16.4 deployments per day.
Release Strategy	Utilizes canary deployments and revisioned APIs through Apigee, allowing for gradual rollouts and minimizing risks.	Employs blue/green and canary deployments via AWS CodeDeploy, enabling safe and incremental releases.	Adopts canary deployments with real-time feedback, leveraging tools like Spinnaker for continuous delivery.	Implements gradual rollouts using feature flags and internal gating mechanisms to ensure stability.

API Architecture	Primarily RESTful APIs managed through Cloud Endpoints and Apigee, supporting scalable and secure API management.	Offers REST and WebSocket APIs via Amazon API Gateway, facilitating flexible and scalable API solutions.	Transitioned to a federated GraphQL architecture using Apollo Federation, promoting modular and scalable API development.	Combines RESTful APIs with an evolving GraphQL layer to provide flexible and efficient API interactions.
API Ownership Model	Centralized platform team oversees APIs, with support for team-managed APIs to encourage collaboration and consistency.	Service teams have full ownership of their APIs, promoting autonomy and rapid development cycles.	Decentralized ownership where each team manages its own GraphQL subgraph, fostering independence and scalability.	Each team is responsible for its API domain, ensuring accountability and streamlined development processes.
CI/CD Tooling	Employs Cloud Build, GitHub Actions, and Cloud Deploy to automate build, test, and deployment processes.	Utilizes AWS CodePipeline, CodeBuild, and CodeDeploy to streamline continuous integration and delivery workflows.	Leverages Spinnaker, Jenkins, and GitHub Actions to facilitate continuous integration and deployment across services.	Operates a custom deployment pipeline integrated with GitHub Actions to support continuous delivery.
Testing Strategy	Incorporates pre-deployment tests, schema validation, and traffic splitting to ensure reliability and performance.	Integrates automated testing using Postman/Newman, along with unit and integration tests to validate functionality.	Implements integration tests and schema checks within CI pipelines, complemented by end-to-end tests for comprehensive coverage.	Conducts extensive integration, unit, and API tests with schema contracts and mocks to maintain high-quality standards.
Automated Schema Checks	Yes; utilizes Apigee and custom validation tools to enforce schema integrity and compatibility.	Partial; employs Postman tests and custom tooling to perform schema validations.	Yes; integrates Apollo Federation validation and CI tools to ensure schema consistency across subgraphs.	Yes; uses custom schema validators and contract testing within CI pipelines to uphold API standards.
Deployment Independence	Partial; while some central governance exists, teams have the flexibility to manage their deployments within defined parameters.	Yes; individual service teams can deploy independently using API Gateway and Lambda, enhancing agility.	Yes; each subgraph is deployed independently, allowing teams to iterate without affecting others.	Yes; teams deploy their respective APIs independently, facilitating rapid and isolated updates.
Monitoring Tools	Employs Apigee Analytics and Cloud Monitoring to track API performance and health metrics.	Uses AWS CloudWatch and X-Ray for comprehensive monitoring and tracing of API activities.	Utilizes Atlas, Netflix's internal monitoring system, along with custom dashboards for observability.	Implements an internal observability stack, including Prometheus and custom logging solutions, to monitor API operations.
Rollback Mechanism	Supports revision-based rollbacks in Apigee, enabling quick reversion to previous stable states.	Facilitates automated rollbacks through CodeDeploy and monitoring alarms to mitigate deployment issues.	Employs Spinnaker's automated rollback capabilities to swiftly address deployment failures.	Utilizes feature flags and CI/CD mechanisms to perform gradual rollbacks and instant reverts when necessary.
API Testing in CI/CD	Mandatory; integrated into Cloud Build pipelines to ensure API reliability before deployment.	Integrated within CodeBuild and pre-deployment stages to validate API functionality and performance.	Fully integrated into CI pipelines, encompassing various testing stages to maintain API quality.	Deeply embedded within CI/CD processes, employing mocks, schema diffing, and simulation environments for thorough testing.



Number of Deployments made by Companies in a day

Why API Testing Needs a Strategic Reboot



Most companies now deploy APIs with agility—faster releases, safe rollouts, automated pipelines. But faster delivery doesn't mean better quality. In fact, some teams are releasing APIs faster without even thinking about the safety index of these tests.

The reason? A quiet bias.

Developer Bias Is Real — and Risky

It starts with good intentions. Developers write unit tests, verify core functionality, and ship. But over time, teams assume these tests are enough. QA steps in late—or not at all. UI testing becomes the default, and APIs go untested in edge cases, integrations, and under load.

This is what we call **developer bias**: the false confidence that developer-written tests fully cover real-world behavior.

Research suggests this can lead to overlooked issues, with Gartner reporting that comprehensive API testing can reduce testing time by **50% and production defects by 30%**.

Reasons for decreased focus include tight deadlines, sole focus on UI testing, over-reliance on automation, and underestimating API vulnerabilities, especially as security threats grow, with **71% of web traffic** now API calls.

Teams should refocus by involving QA teams for independent validation, ensuring comprehensive test coverage beyond developer assumptions. This approach aligns with the global trend toward API-first strategies, where testing is integral to delivering reliable, scalable services.

The Importance of Functional Testing and Breaking the Bias

Functional testing verifies that APIs meet specified requirements, crucial for ensuring reliability. However, some testers may fall into a “**testing bias**,” changing URLs and writing basic test cases that pass QA unchecked, often influenced by developer-provided tests.


This developer bias can lead to gaps in coverage, missing edge cases or error handling, as developers might assume their code is correct.

To break this, QA teams should design test cases independently, focusing on:

- **Edge Cases:** Testing invalid inputs, like missing authentication, to ensure robust error handling.
- **Integration Points:** Verifying how the API interacts with other systems, not just standalone functionality.

- **Exploratory Testing:** Discovering issues beyond scripted tests, challenging developer assumptions.

By doing so, teams ensure APIs are thoroughly validated, reducing risks and enhancing user trust, especially as APIs power critical applications like e-commerce and healthcare.



Why Load Testing, Performance Testing, Validation, and Process Performance Testing Matter

Beyond functionality, APIs must handle real-world demands:

- **Load Testing:** Ensures the API can manage high traffic, critical for scalability during peak usage, like Amazon’s Black Friday sales ([Testing phase - Development and Test on Amazon Web Services](#))
- **Performance Testing:** With qAPI, teams can simulate realistic user loads, test for SLA compliance, and integrate traffic modeling into their QA routines—**without needing a separate performance environment.**

- **Validation:** confirm data integrity and adherence to API contracts, preventing downstream errors, especially in financial APIs.
- **Process Performance Testing:** Evaluates end-to-end workflows, ensuring seamless integration, vital for microservices architectures.

These testing types collectively ensure APIs are reliable, performant, and secure, aligning with the global push for high-quality digital services.

To refocus, teams should involve QA teams for independent validation, ensuring comprehensive test coverage. This aligns with global trends, where API-first strategies are becoming standard,

and testing is integral to delivering reliable services. Practical steps include:

- **Training Programs:** Educate teams on API testing importance, using tools like qAPI for hands-on practice.
- **QA Empowerment:** Ensure QA teams have autonomy to design test cases, challenging developer assumptions.
- **Security Integration:** Incorporate security testing into API workflows, addressing the growing threat landscape.

By refocusing, teams can mitigate risks and enhance API quality, crucial for maintaining customer trust and operational efficiency.

Functional Testing That Goes Beyond “Happy Paths”

Traditional QA approaches sometimes reinforce the same biases they aim to catch. A tester receives a spec, copies a few sample requests, runs them, and signs off. But this only validates that the API works under normal conditions—not that it can handle exceptions, edge cases, or unexpected behavior.

To break that loop, induce chaos testing by:

- **Testing for the unexpected:** missing tokens, invalid formats, incorrect HTTP methods.

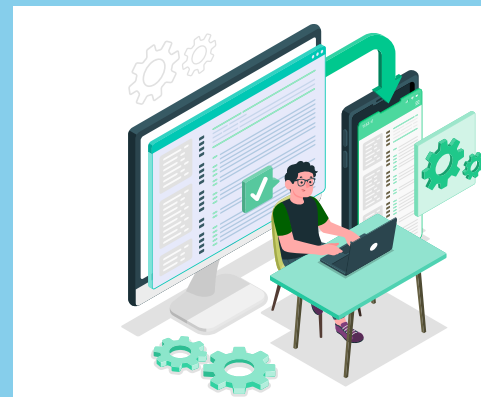
- **Simulating integration:** how does the API behave when a downstream service is slow or unavailable?
- **Validating responses:** are error codes accurate? Are response bodies consistent with OpenAPI contracts?

Platforms like **qAPI** actually work on embedding schema validation and automated regression checks into every CI/CD pipeline, ensuring functional gaps don't reach production undetected.

To break this, QA teams should design test cases independently, focusing on:

- **Edge Cases:** Testing invalid inputs, like missing authentication, to ensure robust error handling, crucial for APIs powering critical applications like healthcare.
- **Integration Points:** Verifying how the API interacts with other systems, not just standalone functionality, aligning with microservices trends.
- **Exploratory Testing:** Discovering issues beyond scripted tests, challenging developer assumptions, enhancing coverage for complex APIs.


By adopting this methodology, teams ensure APIs are thoroughly validated, reducing risks and enhancing user trust, especially as APIs power critical applications like e-commerce and financial services.



Conclusion: The Future of API Testing

As APIs continue to be the foundation of modern software ecosystems, the importance of robust testing practices cannot be overstated. Emerging trends, such as AI-driven testing and continuous integration, will shape the future, with **40% of DevOps teams** expected to use AI-driven tools by 2027.

For organizations to stay competitive, investing in comprehensive API testing is essential, ensuring reliability, performance, and security, thereby driving innovation and customer satisfaction in the digital age.





About Us

qAPI, part of Qyrus, is a leading codeless API testing platform that specializes in delivering advanced cloudbased testing solutions. We help businesses with innovative tools and services designed to streamline API testing, ensure reliability, and enhance application performance. Trusted by financial institutions, logistics companies, and many more worldwide, we help organizations create products and APIs they can depend on for seamless performance and integration.

To learn more about our products and services, visit us at qyrus.com/qapi



Contact Us

+004 02154 15859

Qyrus INC, 1 S Wacker Dr, Suite 3150,

Chicago, IL, 60606

USA

www.qyrus.com/qapi