



Heroku 101

2012

2 min read • 479 words

Introduction

This presentation provided a comprehensive introduction to Heroku's platform-as-a-service offering, focusing on how its elegant abstractions dramatically simplify application deployment and management for developers who want to focus on building features rather than managing infrastructure.

Key Concepts

The presentation began by clarifying the often-confusing landscape of cloud service models. **Software as a Service (SaaS)** serves software users by providing feature-rich applications with transparent updates—platforms like Facebook and Trello exemplify this approach. **Infrastructure as a Service (IaaS)** targets operations teams, offering on-demand machine resources through services like AWS and Digital Ocean.

Platform as a Service (PaaS) occupies the sweet spot for developers, providing transparent updates and managed infrastructure without requiring server management expertise.

PaaS represented a fundamental shift in developer workflow, abstracting away infrastructure concerns and enabling rapid prototyping and deployment that influenced the entire industry.

Heroku and Google App Engine pioneered this model, demonstrating how developers could deploy applications without becoming system administrators.

Using Heroku with Python

The Python deployment workflow demonstrates Heroku's elegance through its simplicity. **Locally**, developers follow familiar patterns: create a virtual environment, install dependencies using `pip`, and run a development server with standard commands like `python manage.py runserver 0.0.0.0:8000`.

On Heroku, this local workflow translates seamlessly to production. Dependencies are still managed through `pip`, but deployment is orchestrated through a `Procfile`—a simple text file that specifies exactly which commands Heroku should execute to start your application. This approach maintains consistency between development and production environments while abstracting away the complexities of server configuration.

Understanding Dynos

The **dyno** represents Heroku's fundamental abstraction—a lightweight process running in a container rather than a traditional server or virtual machine.

Heroku's dyno model was an early implementation of containerization concepts that later evolved into Docker and Kubernetes orchestration patterns.

This abstraction enables effortless horizontal scaling through simple commands like `heroku scale`, allowing applications to handle varying load patterns without requiring deep infrastructure knowledge. Dynos can be allocated to different process types, enabling sophisticated multi-process web applications with worker processes, scheduled tasks, and web servers all managed through the same unified interface.

Managing Applications

Heroku's application management capabilities extend far beyond simple deployment. The `heroku run` command enables **arbitrary command execution** within the application environment—whether running database migrations,

opening interactive Python shells, or executing maintenance scripts. This capability maintains the development experience while operating in production infrastructure.

Addon services represent one of Heroku's most powerful features, providing managed infrastructure resources like PostgreSQL, Redis, and Kafka through a unified interface.

The addon ecosystem pioneered the "marketplace" model for cloud services, allowing third-party providers to integrate seamlessly with the platform through standardized APIs.

These services integrate seamlessly with applications through environment variables and standardized connection patterns, eliminating the complexity of manual service provisioning and configuration.

Application configuration is managed through environment variables using `heroku config:set`, enabling clean separation between code and configuration while supporting environment-specific settings across development, staging, and production deployments.