

Lecture 4: Containers II

AC215

Pavlos Protopapas
SEAS/Harvard



Logistics/Reminders

- If you have formed groups - please update [group info sheet](#)
- Please fill out survey -
<https://canvas.harvard.edu/courses/136127/assignments/866239>
(Survey responses have been updated)
- Office Hours details here -
<https://edstem.org/us/courses/58478/discussion/5229430>

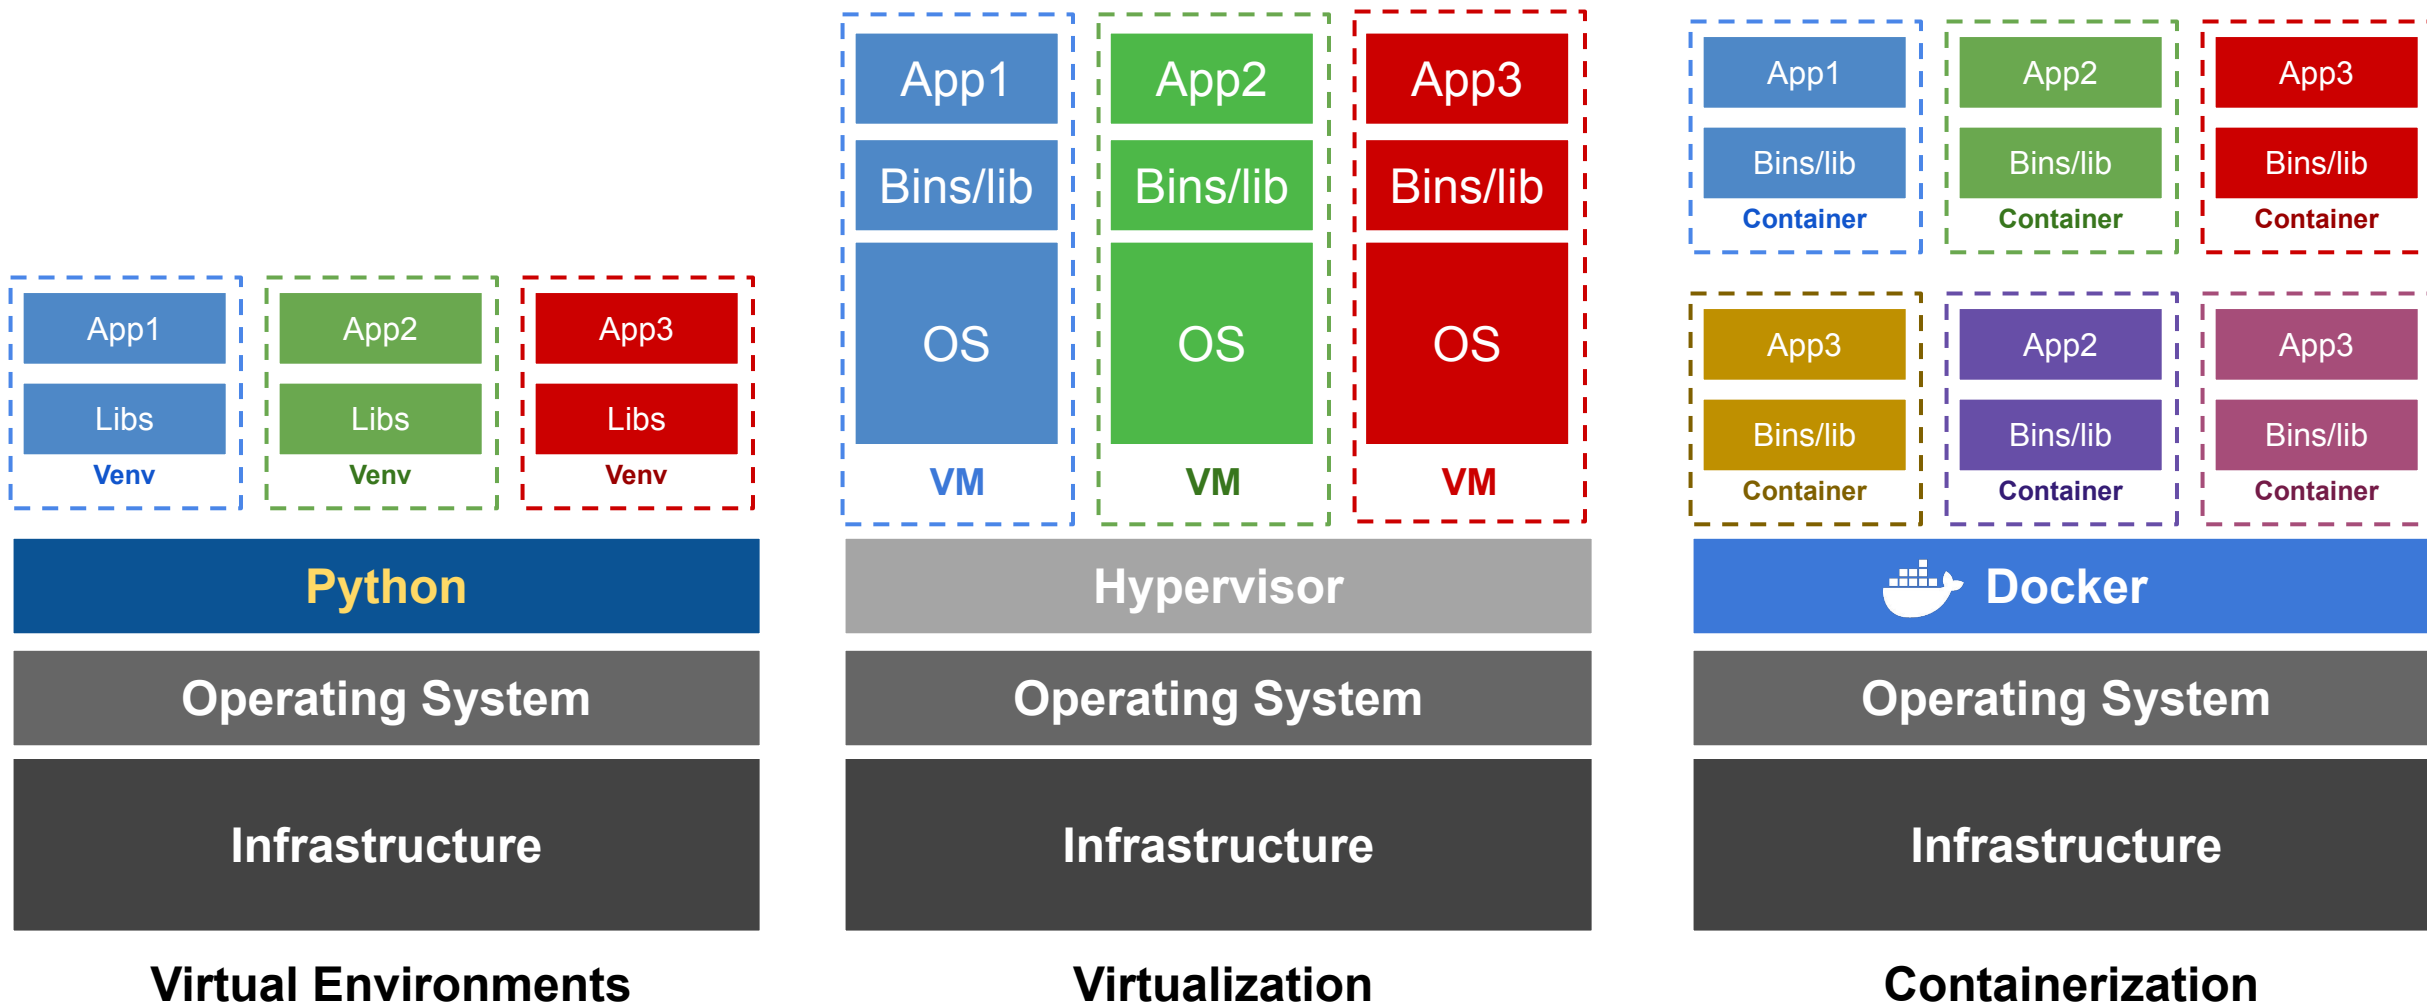
Outline

1. Recap: Review of Previous Material
2. Containers in Architecture: Microservices vs. Monolithic
3. Implementing Containers as Microservices

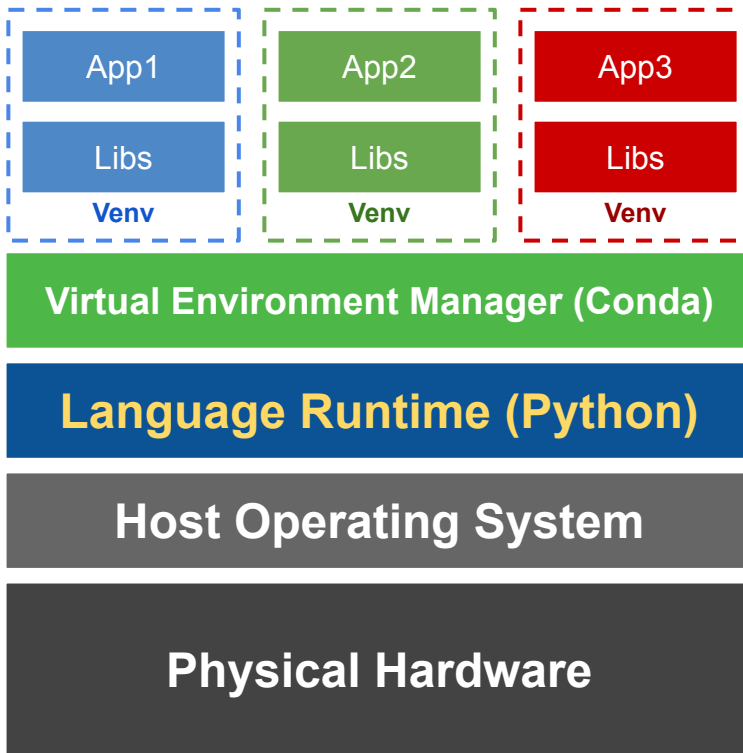
Outline

1. **Recap: Review of Previous Material**
2. Containers in Architecture: Microservices vs. Monolithic
3. Implementing Containers as Microservices

Recap: Environments vs Virtualization vs Containerization



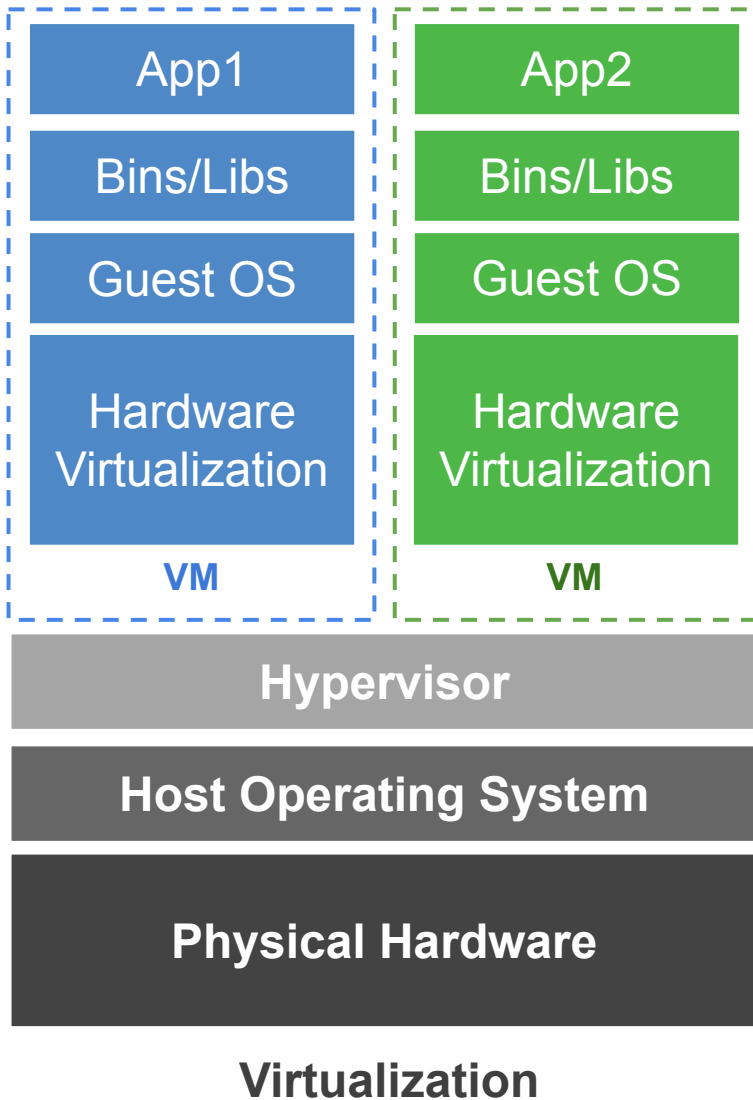
Environments



Virtual Environments

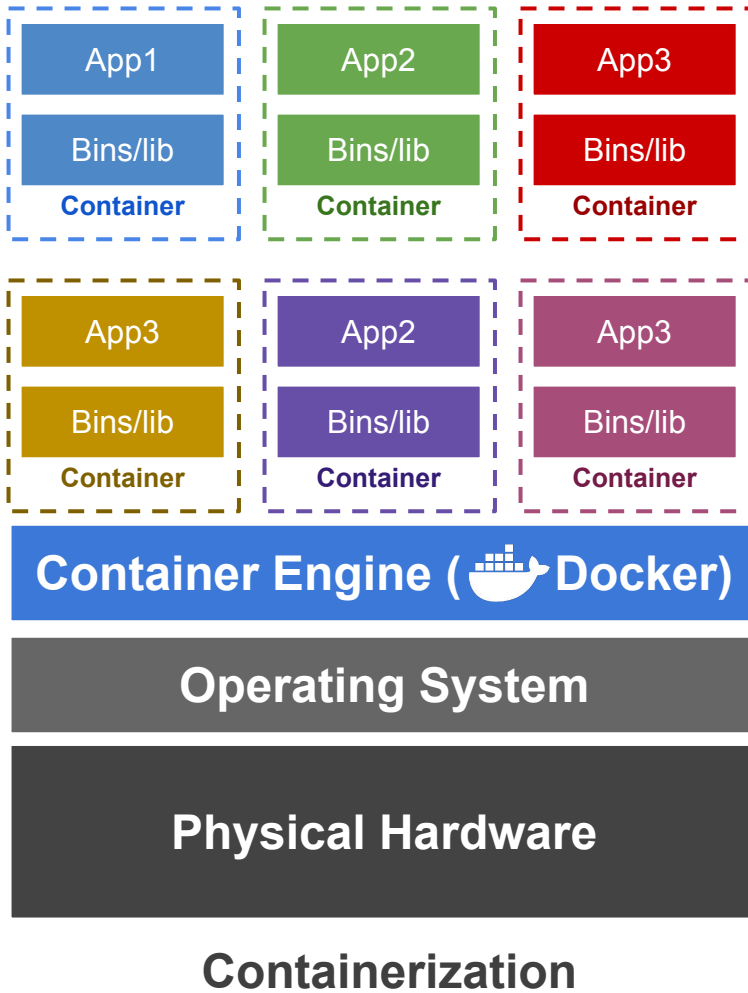
- **Dependency Isolation:** Virtual environments redirect dependencies to their own directories, avoiding system-wide installs.
- **No Kernel Isolation:** Unlike VMs or containers, they don't isolate the kernel.
- **Resource Efficiency:** Without an OS or kernel, virtual environments are lightweight and resource-efficient.
- **Filesystem Access:** Files written within a virtual environment can be accessed from other environments, as there's no filesystem isolation.

Virtualization (Virtual Machines)



- **CPU Virtualization:** Virtual CPUs are mapped to physical cores, but hypervisor management adds some overhead.
- **Emulated Devices:** VMs use virtual devices (CPUs, network adapters, disks) translated by the hypervisor to real hardware.
- **Full OS:** Each VM runs its own guest OS with independent kernel and user spaces, but this reduces efficiency.
- **Resource Allocation:** RAM, CPU, and disk space are often allocated in fixed blocks, limiting flexibility in resource usage.

Containerization



- **Namespaces:** Containers isolate processes and resources, making them act like independent systems. For example, PID namespaces separate process IDs, and mount namespaces provide unique file systems.
- **Cgroups:** These limit CPU, memory, and IO usage for each container, ensuring efficient resource use.
- **Process Virtualization:** Namespaces and cgroups work together to isolate and control processes.
- **Shared Kernel:** Containers use the host's OS kernel but have their own files, making them lightweight and efficient.
- **Direct Access:** Containers interact with host resources directly, reducing overhead compared to VMs.

Outline

1. Recap: Review of Previous Material
2. **Containers in Architecture: Microservices vs. Monolithic**
3. Implementing Containers as Microservices

Why use Containers?

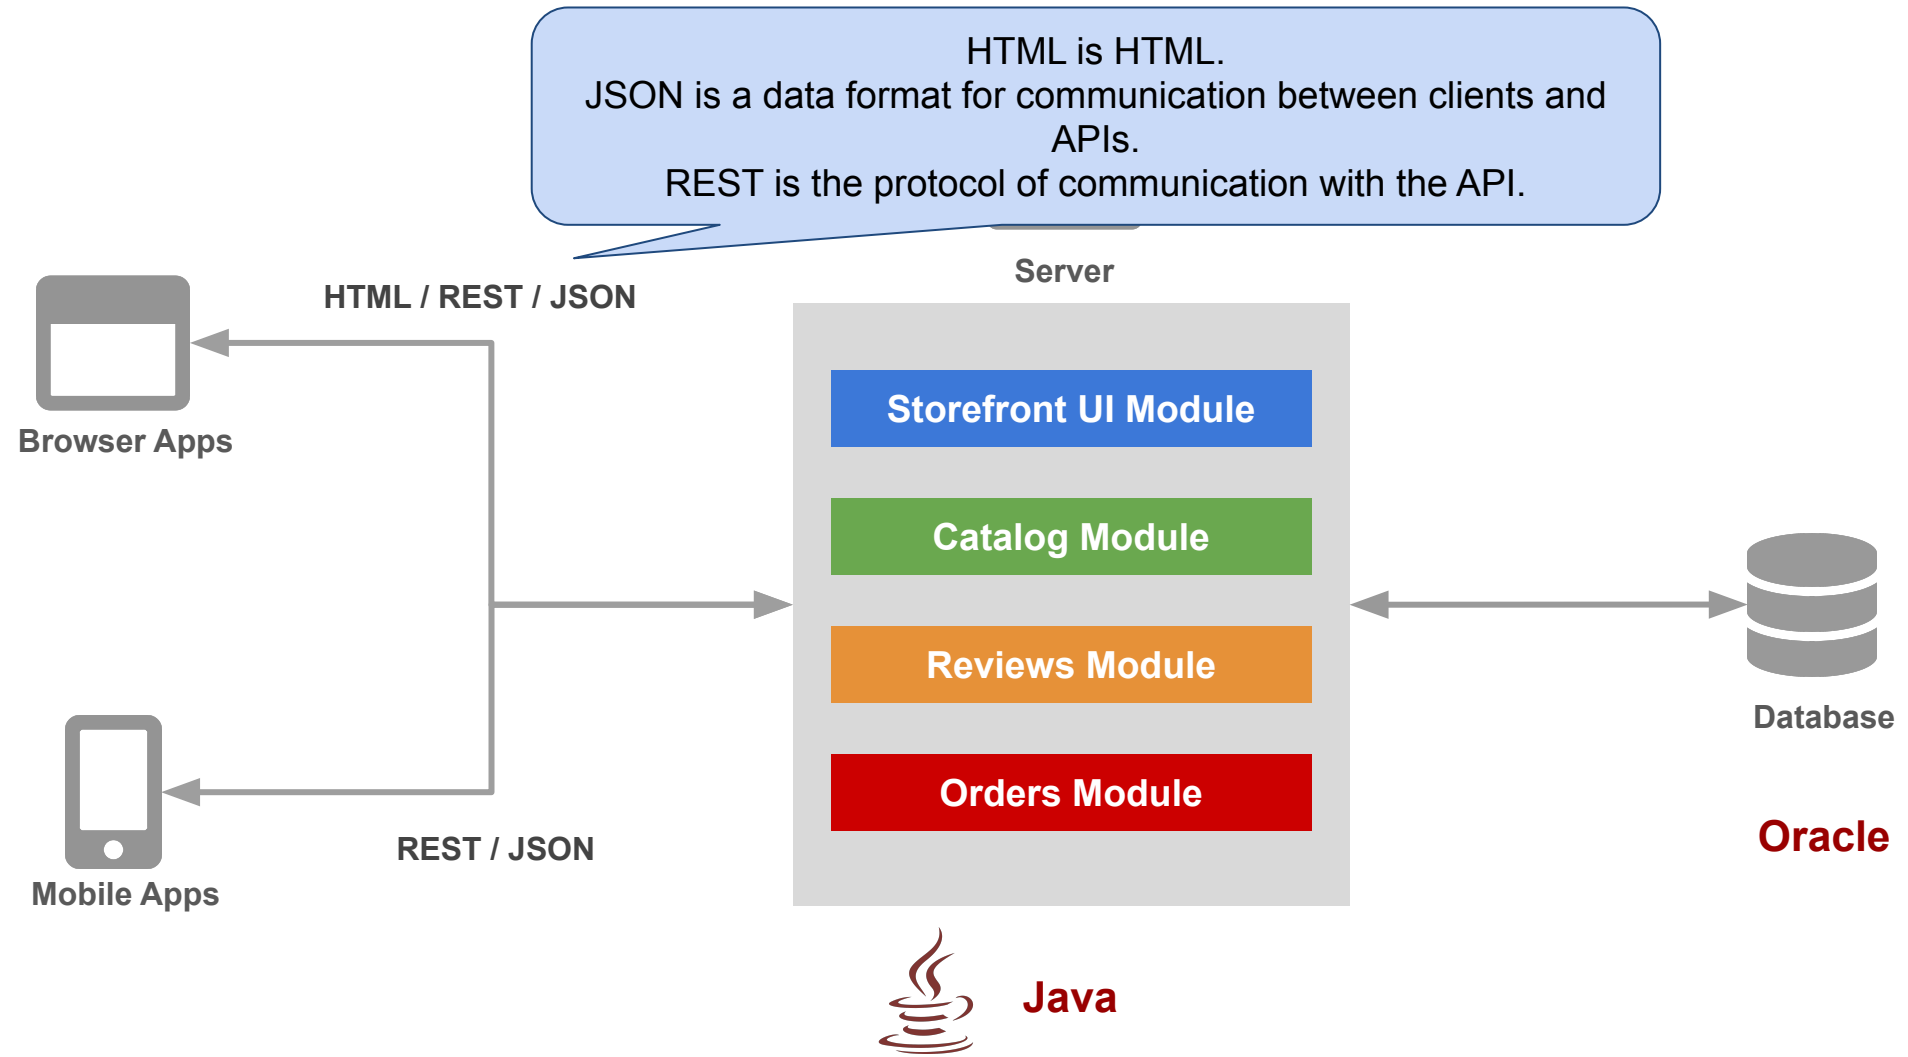
Conceptual Scenario

- Picture building an application, such as an online cheese store.

Traditional Approach

- Traditionally you would build this using a [Monolithic Architecture](#).

Monolithic Architecture



Monolithic Architecture - **Advantages**

Simplicity in Development:

Most tools and IDEs natively support monolithic applications.

Ease of Deployment:

All components bundled into a single, unified package.

Scalability:

Easier to scale by replicating the entire application as a whole (horizontal scaling).

Monolithic Architecture - Disadvantages

Maintenance Challenges:

Complexity increases over time, making it harder to implement changes or find issues.

System Vulnerability:

A failure in a single component can lead to the collapse of the entire system.

Patching Difficulties:

Patching or updating specific modules can be cumbersome due to tightly-coupled components.

Monolithic Architecture - Disadvantages

Technology Lock-in:

Adopting new technologies or updating existing ones can be problematic due to interdependencies.

Slow Startup:

Increased startup time as all components must be initialized simultaneously.

Onboarding Challenges:

New users need to familiarize themselves with the entire codebase.

Applications have changed dramatically

A decade ago

Apps were monolithic
Built on a single stack (e.g. .Net or Java)
Long lived
Deployed to a single server

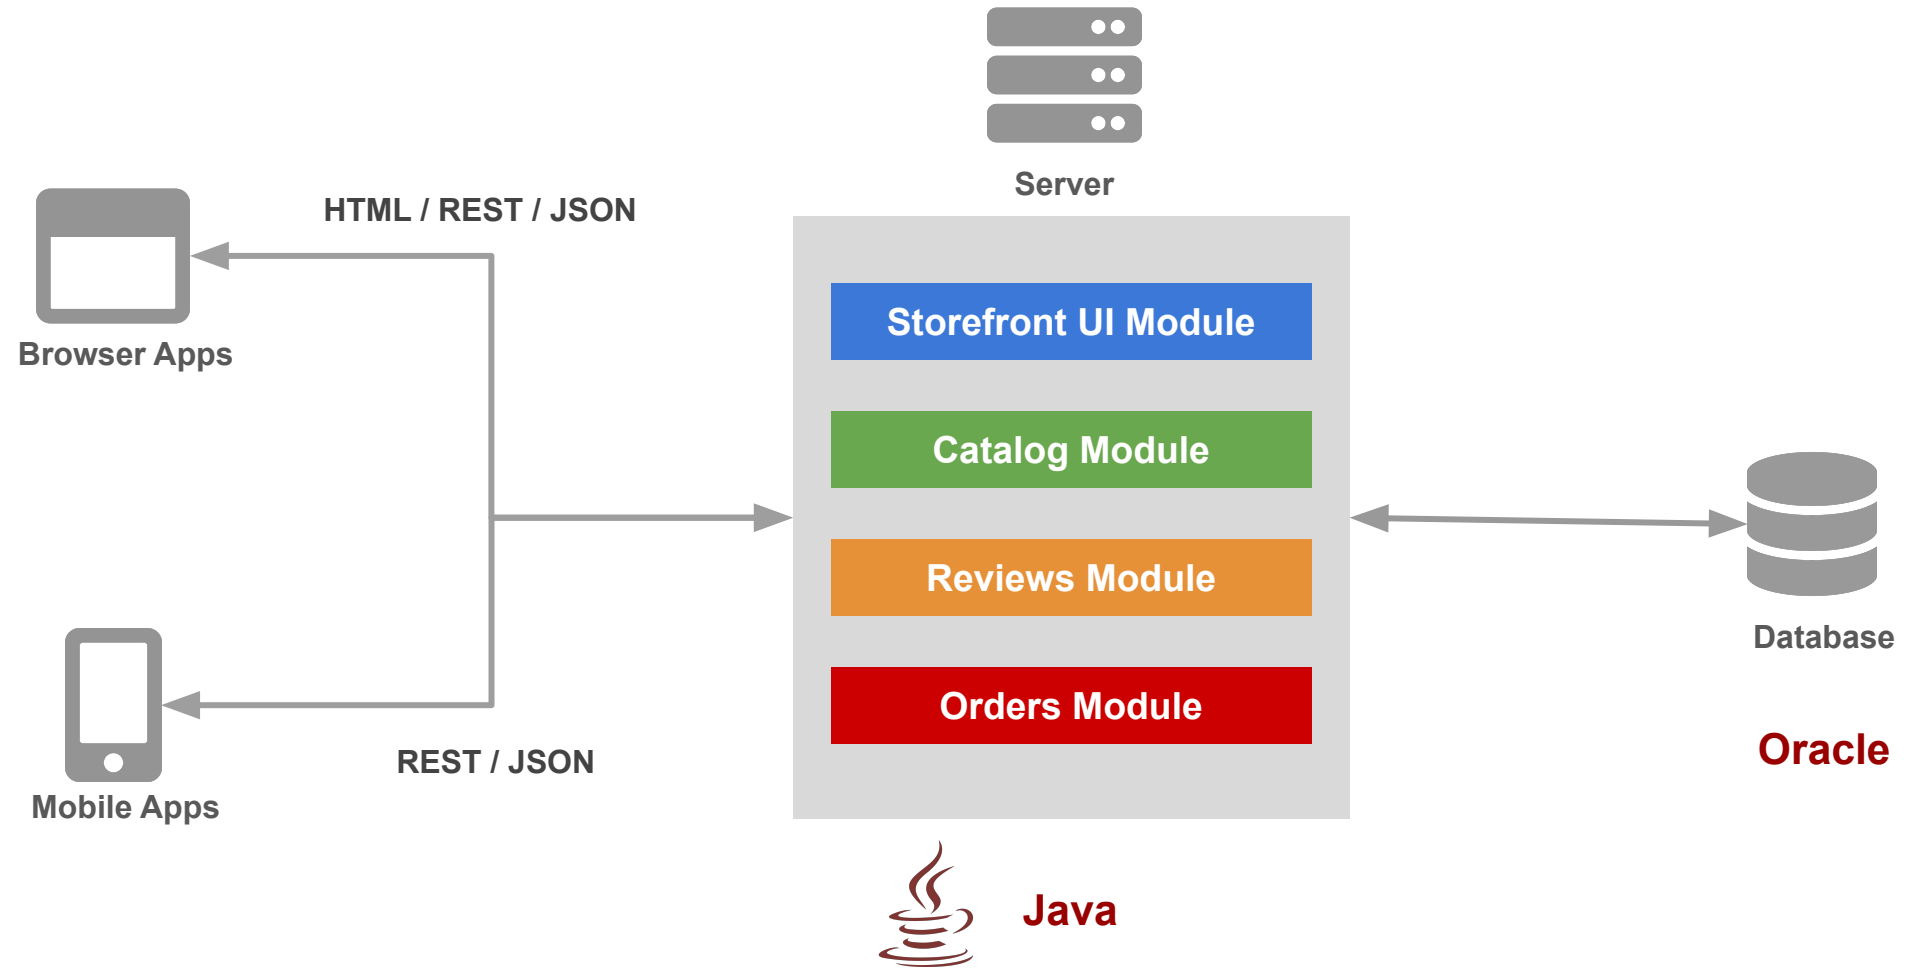
Today

Apps are constantly being developed
Build from loosely coupled components
Newer version are deployed often
Deployed to a multitude of servers

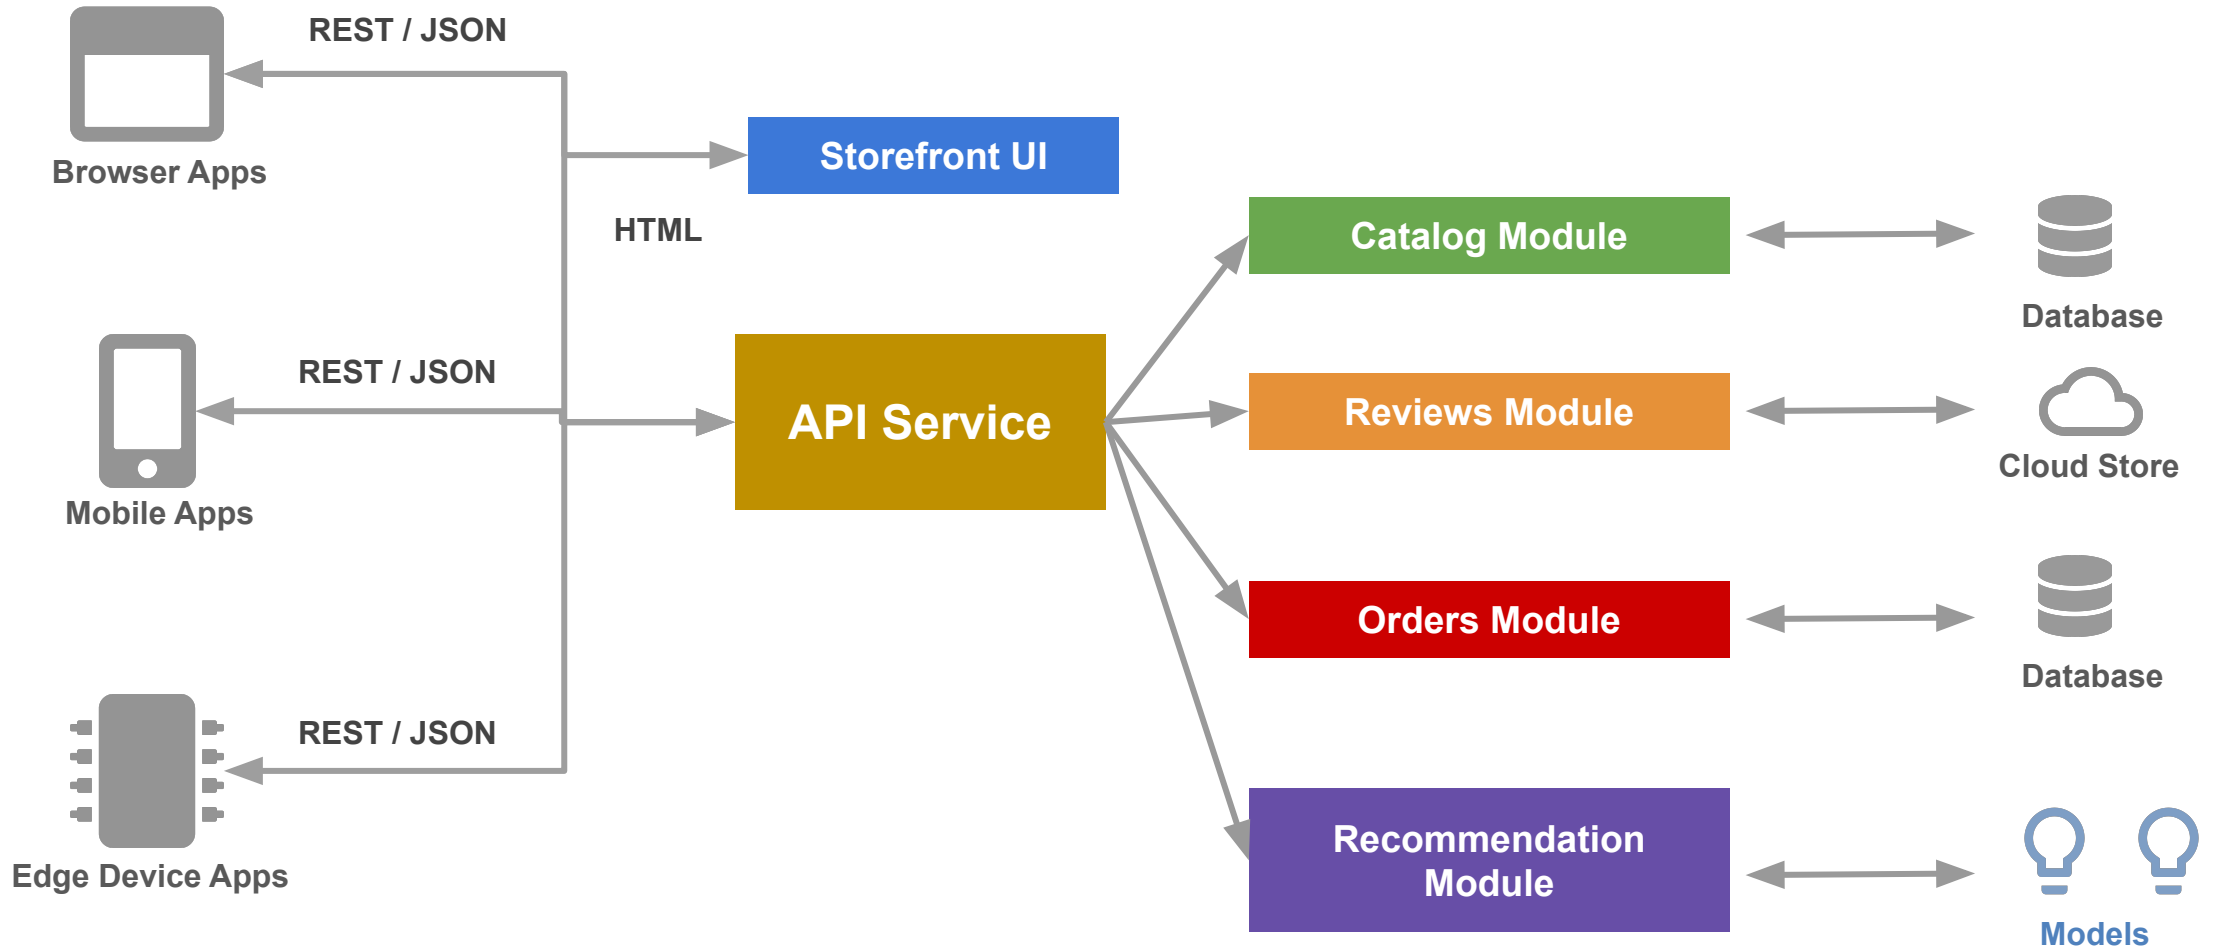
Data Science

Apps are being integrated with various data types/sources and models

Monolithic Architecture



Today: Microservice Architecture



Microservice Architecture - **Advantages**

Simplified Maintenance:

Modular design makes it easier to manage, update, and debug individual services.

Fault Isolation:

Independent components ensure that failure in one service doesn't bring down the entire application.

Streamlined Patching:

Easier to patch or update specific services without affecting the entire system.

Technological Flexibility:

Adapting to or adopting new technologies becomes seamless due to service independence.

Quick Startup:

Reduced startup time as all components can be initialized in parallel.

Microservice Architecture - Disadvantages

Development and Deployment Complexity:

Using multiple technologies across components can complicate both development and deployment, as managing diverse dependencies requires a more intricate setup.

Scaling Concerns:


Scaling the entire application can be intricate due to disparate components.

Docker + Kubernetes

Why use Containers?

- Consider a software development team workflow for developing an App
- Traditionality you would develop/build this independently in various machines (dev, test, qa, prod)

Software Development Workflow (no Docker)



OS Specific **installation** in every developer machine


Node.js is a platform that allows you to run JavaScript code outside a web browser.

Source Control



Every team member moves code to source control


Build Server



The build server must be configured with all **necessary software** and **frameworks**.

The production build process involves pulling the code from the source control system.


Test Servers / Production



The production server must have all **necessary** software and frameworks **installed**.

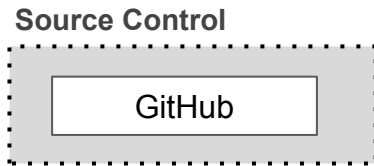
It will also run on a different OS version compared to the development machines.

Software Development Workflow (with Docker)

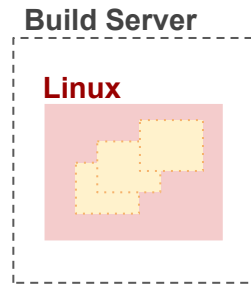


Development machines only needs **Docker installed**.

Containers need to be setup only once.

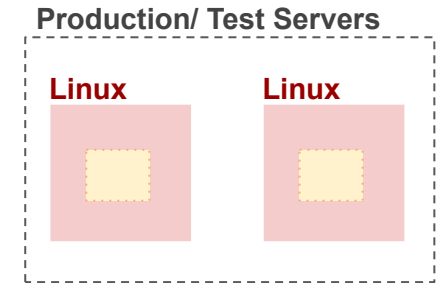


Every team member moves code to source .



Build server only needs **Docker installed**.

Docker **images** are built for a release and pushed to **container registry**.



Production server only needs **Docker installed**.

Production server pulls Docker **images** from **container registry** and runs them.


Software Development Workflow (with Docker)

Who creates the Dockerfile, and where is it stored? Do we use pre-built images or does each developer build them? Who is in charge of managing this? Also, what's the process for handling the Pipfile and Pipfile.lock?

Development machines only needs **Docker installed**.
Containers need to be setup only once.

This seems like a lot.

Software Development Workflow (with Docker)

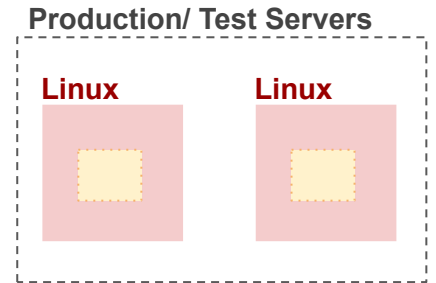



Development machines only needs **Docker installed**.

Containers need to be setup only once.



Software
Every code



Production server only needs **Docker installed**.

Production server pulls Docker **images** from **container registry** and runs them.

Outline

1. Recap: Review of Previous Material
2. Containers in Architecture: Microservices vs. Monolithic
- 3. Implementing Containers as Microservices**

Tutorial (T5) - Building the Mega Pipeline App

One of [Formaggio.me](https://formaggio.me)'s goals is to create a podcast on various cheese-related topics. After recording the podcast, we plan to transcribe the audio, use a language model to correct grammar and enhance the text, and then generate audio that will be made available to our users. Remember, we aim to reach an audience all over the world, so the podcast will be translated into various languages and synthesized into audio.

The goal here is to simulate a realistic development scenario where each component will be developed by different teams and containerized.

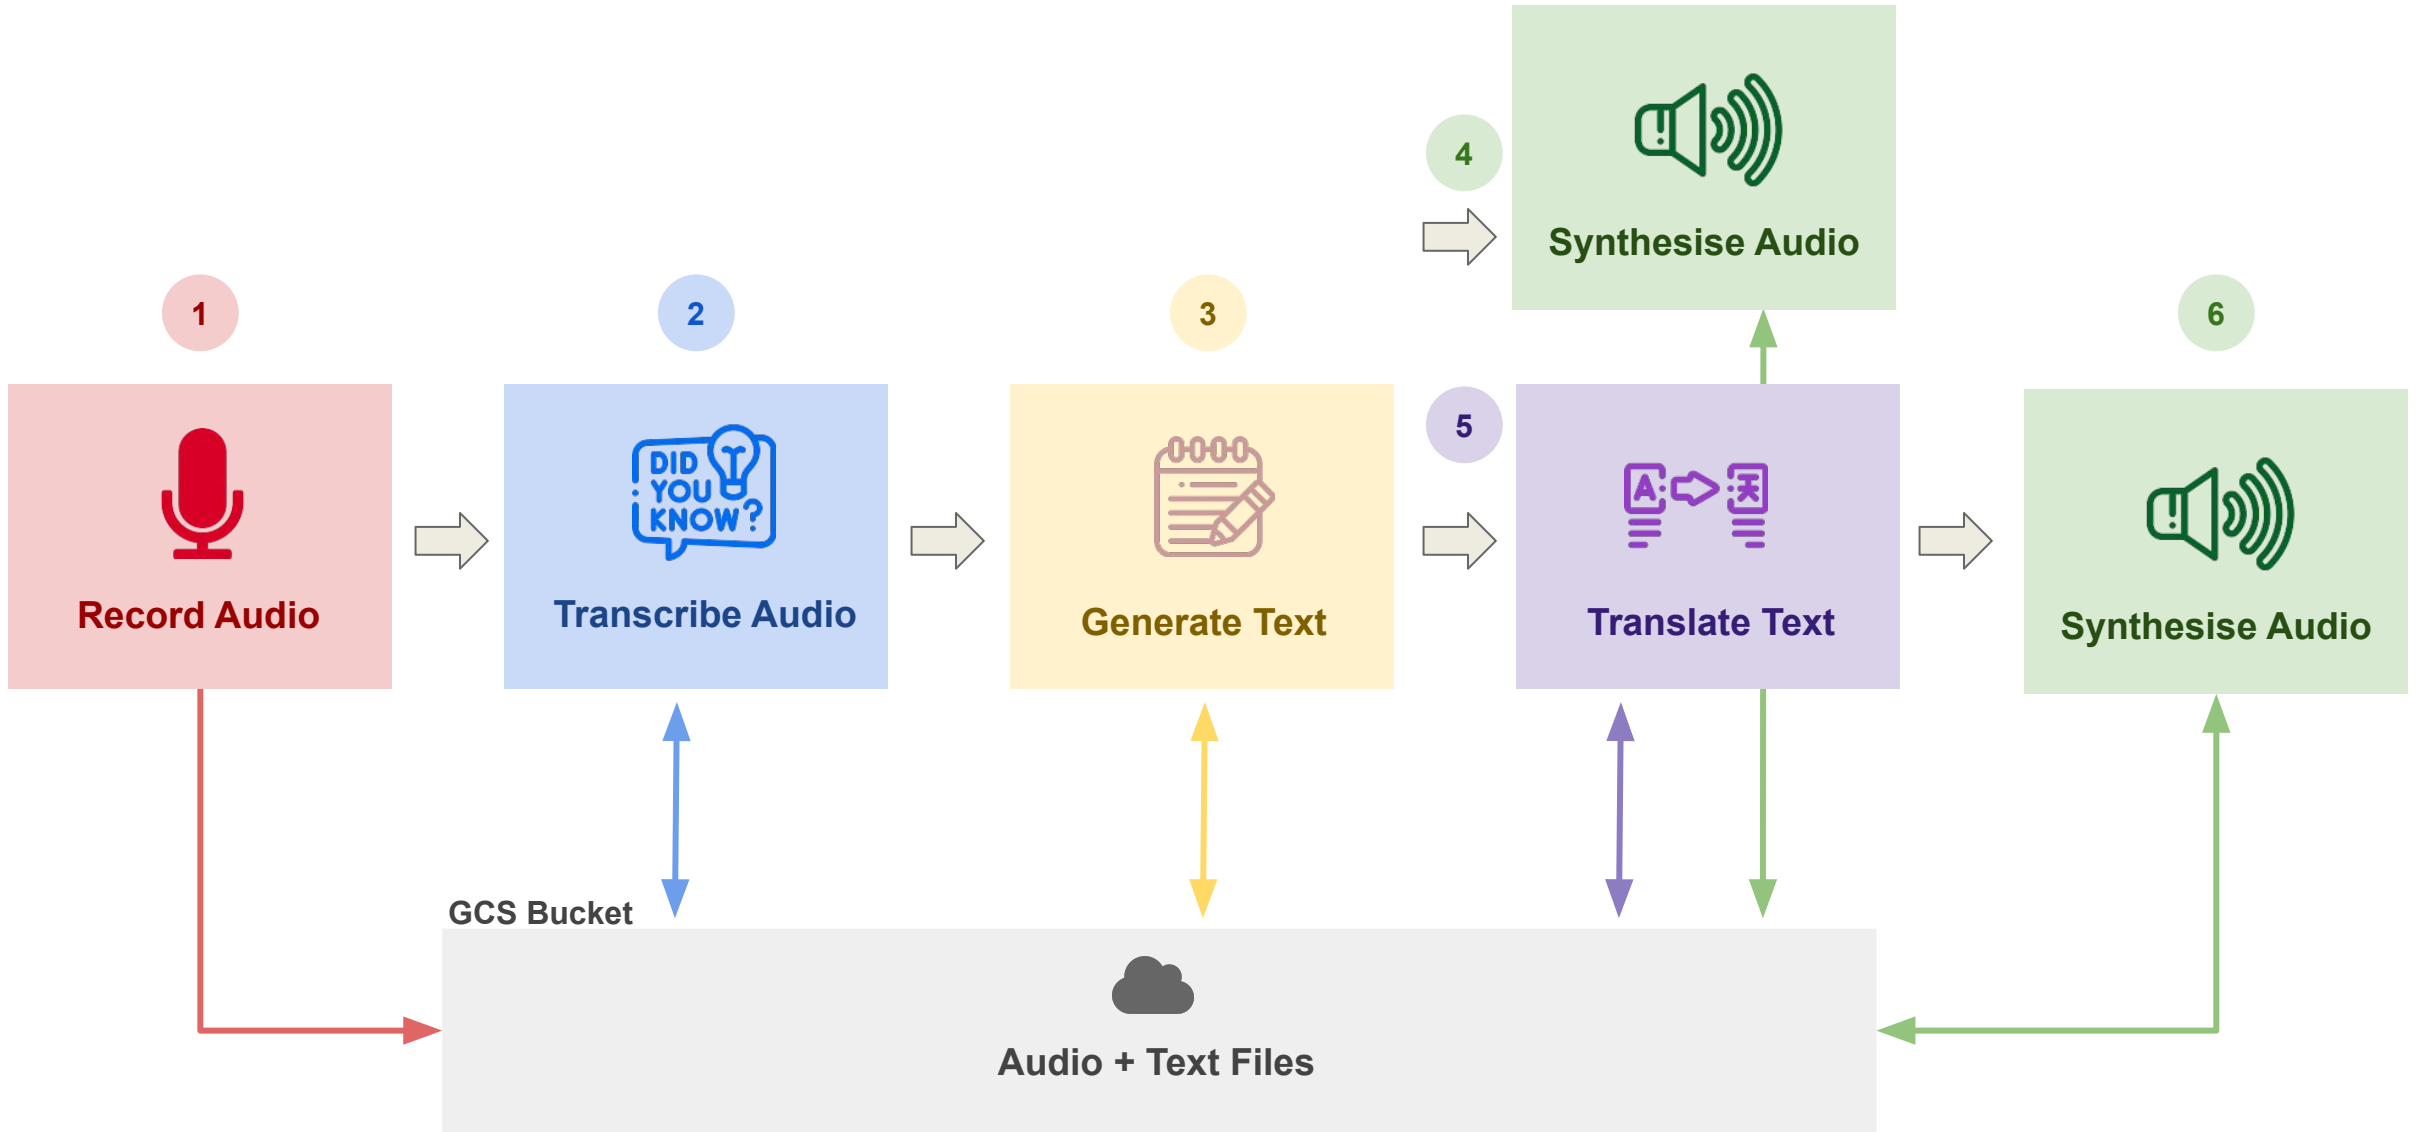
BONUS: You can use elevenlabs API to generate text with Pavlos' voice or your own voice.



Tutorial (T5) - Building the Mega Pipeline App

1. Role of the bucket and that they will be using our bucket. They only need the secrets to communicate
2. Remove all unnecessary folders
3. Sequence matters
4. All containers are local
5. Basic steps are: Create
 - a. Docker, Pipfile, Pipfile.loc and secrets
 - b. Build images
 - c. Run image
 - d. Run cli.py
 - e. Clarify some of the cli.py synthesize options
6. go through process step by step

Tutorial (T5) - Building the Mega Pipeline App




Tutorial (T5) - Team Challenge!

- We'll form teams of 5, and each team must complete 5 tasks.
- The first team to finish all tasks will win a special prize!
- Don't worry, every team that completes the tasks will also get a reward!
- The rewards are a surprise—so give it your best effort and have fun!



Tutorial (T5) - Building the Mega Pipeline App

☰ AC215: Mega Pipeline App

Audio Prompts Transcribed Audio Generated Text Synthesised Audio Translated Text Synthesised Audio 

▶ 0:00 — 🔊

Team: AGSRZ
welcome to the very cheesy podcast where we discuss all things cheese this is your house spotless protopapas we all love c... [show more](#)

Team: group-14a
welcome to the very cheesy podcast where we discuss all things cheese this is your house spotless protopapas we all love c... [show more](#)

Team: group-06
welcome to the very cheesy podcast where we discuss all things cheese this is your house spotless protopapas we all love c... [show more](#)

Team: group-11
welcome to the very cheesy podcast where we discuss all things cheese this is your house spotless protopapas we all love c... [show more](#)

Team: AGSRZ
Welcome to the Very Cheesy Podcast, where we discuss all things cheese. This is your host, Pavlos Protopapas. We all love c... [show more](#)

Team: group-06
Welcome to the Very Cheesy Podcast, where we discuss all things cheese. This is your host, Pavlos Protopapas. We all love c... [show more](#)

Team: group-11
Welcome to the Very Cheesy Podcast, where we discuss all things cheese. This is your host, Pavlos Protopapas. We all love c... [show more](#)

Team: group-01
Welcome to the Very Cheesy Podcast, where we discuss all things cheese. This is your host, Pavlos Protopapas. We all love c... [show more](#)

Team:group-11
▶ 0:00 — 🔊

Team:group-01
▶ 0:00 — 🔊

Team:group10
▶ 0:00 — 🔊

Team:group-08
▶ 0:00 — 🔊

Team:group-15
▶ 0:00 — 🔊

Team:pavlos
▶ 0:00 — 🔊

Team: AGSRZ
Bienvenue dans le podcast très ringard, où nous discutons de tout ce qui concerne le fromage.Ceci est votre hôte, Pav... [show more](#)

Team: group-11
Bienvenue dans le podcast très ringard, où nous discutons de tout ce qui concerne le fromage.Ceci est votre hôte, Pav... [show more](#)

Team: group-01
Bienvenue dans le podcast très ringard, où nous discutons de tout ce qui concerne le fromage.Ceci est votre hôte, Pav... [show more](#)

Team:group-11
▶ 0:00 — 🔊

Team:group-01
▶ 0:00 — 🔊

Team:group10
▶ 0:00 — 🔊

Team:group-08
▶ 0:00 — 🔊

Team:pavlos
▶ 0:00 — 🔊

Team:group-01
▶ 0:00 — 🔊

Team:pavlos
▶ 0:00 — 🔊

GCP Authentication Methods

Before we start let us review how do we authenticate to different services/accounts and APIs

OAuth 2.0:

For user-driven authentication and access.

Service Account (part of IAM in GCP):

For server-to-server interactions requiring automation and high control without user intervention.

API Key:

For lightweight, less secure access to APIs, use cautiously.

Default Service Accounts (part of IAM in GCP):

For Compute Engine, Kubernetes Engine, and App Engine with predefined permissions.

Workload Identity Federation:

For external identities to access GCP securely.

Tutorial (T5) - Building the Mega Pipeline App

- App: <https://ac215-mega-pipeline.dlops.io/>
- Teams
 - 📝 Task A [transcribe_audio](#):
 - 📄 Task B [generate_text](#):
 - 🎧 Task C [synthesis_audio_en](#):
 - 🇫🇷 Task D [translate_text](#):
 - 🎧 Task E [synthesis_audio](#):
- Instructions: <https://github.com/dlops-io/mega-pipeline>







Now check: <https://formaggio.me/>