



DISTRIBUTED SYSTEMS
CS6421
**CLOUD APPLICATION
ARCHITECTURES**

Prof. Tim Wood and Prof. Roozbeh Haghazadeh

LAST TIME...

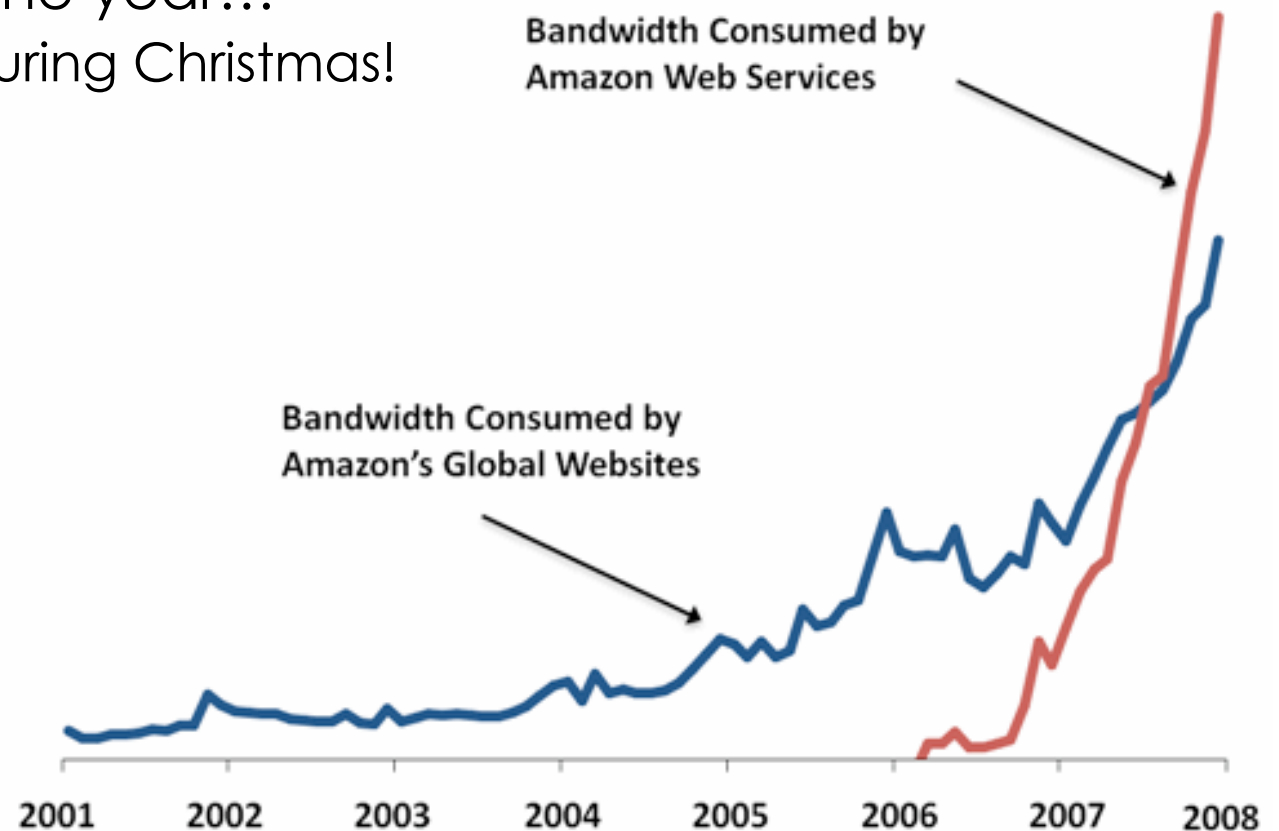
- Performance in Dist. Systems
 - Introduction
 - Performance metrics
 - Models

THIS TIME...

- Cloud Application Architectures
 - IaaS vs PaaS
 - Multi-Tier
 - Microservices
 - Serverless
 - Big Data
 - Machine Learning

AMAZON'S CLOUD

- Amazon built its cloud platform so that other people could pay for its infrastructure during the rest of the year...
 - Only needed peak capacity during Christmas!
- Now its cloud users are far bigger than its own sites



TYPES OF CLOUD SERVICES

- Infrastructure as a Service (**IaaS**)
 - Rent VMs, Containers, physical servers, disks, etc. by the hour
 - Examples: EC2, EBS, S3

 - Benefits?
 - Limitations?
- Other options: Function as a Service, Software as a Service, Resource as a Service...
- Platform as a Service (**PaaS**)
 - Cloud provides a software layer on top of its resources
 - Exposes a programming API for users to develop cloud-based apps
 - Cloud provider manages all underlying resources (autoscaling)
 - Examples: Beanstalk, Lambda, EMR


 - Benefits?
 - Limitations?

MULTI-TIER APPLICATIONS



MULTI-TIER WEB APPLICATIONS

- Traditionally composed of 3 components
- Separation of duties:
 - **Front-end web server** for static content (Apache, lighttpd, nginx)
 - **Application (API) tier** for dynamic logic (PHP, Tomcat, node.js)
 - **Database back-end** holds state (MySQL, MongoDB, Postgres)
- Why divide up in this way?

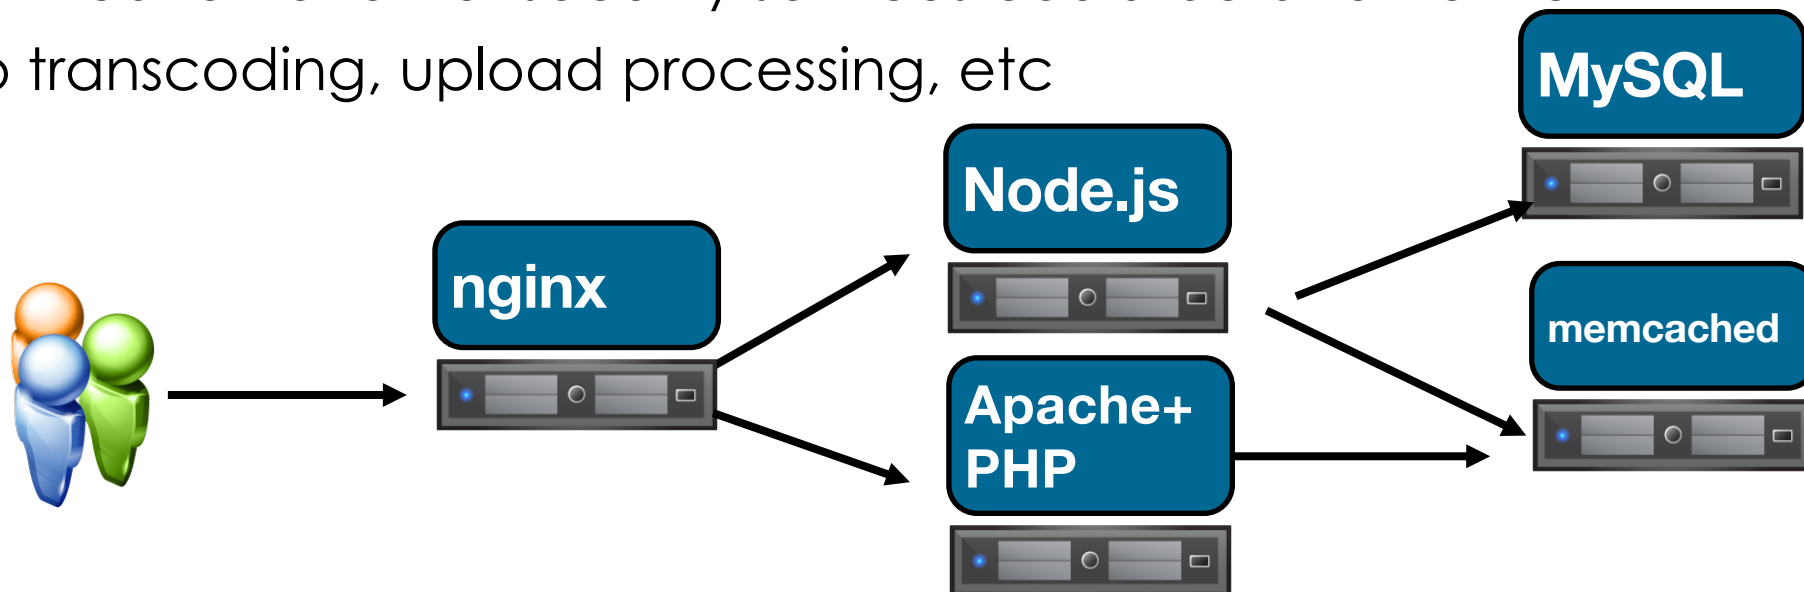


STATEFUL VS STATELESS

- The multi-tier architecture is based largely around whether a tier needs to worry about state
- Front-end - totally **stateless**
 - There is no data that must be maintained by the server to handle subsequent requests
- Application tier - maintains **per-connection state**
 - There is some temporary data related to each user, e.g., my shopping cart
 - May not be critical for reliability - might just store in memory
- Database tier - global state
 - Maintains the global data that application tier might need
 - Persists state and ensures it is consistent


N-TIER WEB APPLICATIONS

- Sometimes 3 tiers isn't quite right
- Database is often a bottleneck
 - Add a cache! (stateful, but not persistent)
- Authentication or other security services could be another tier
- Video transcoding, upload processing, etc



REPLICATED N-TIER

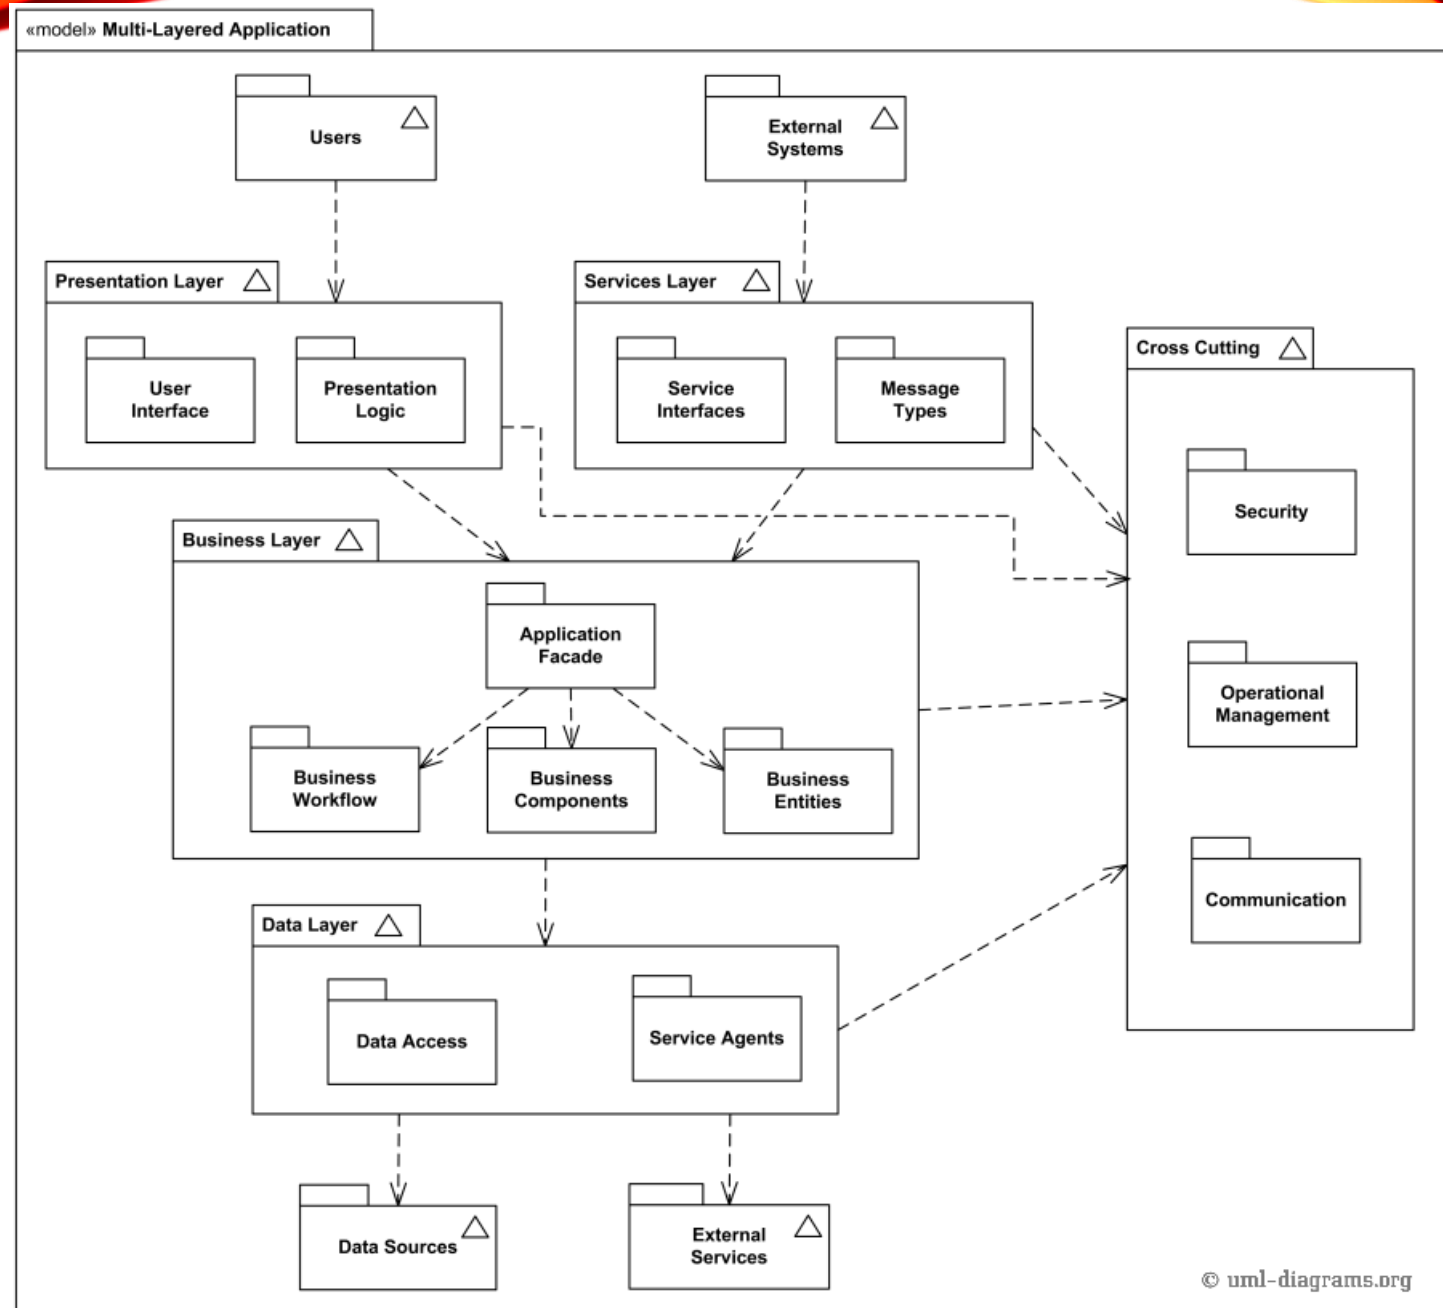
- Replicate the portions of the system that are likely to become overloaded
- How easy to scale...?
 - Apache serving static content
 - Tomcat Java application managing user shopping carts
 - MySQL cluster storing products and completed orders



Tune number of replicas based on demand at each tier

MULTI-LAYERED APPLICATION

An example of UML model diagram representing a model of a layered application, based on the Microsoft Application Architecture Guide, 2nd Ed.

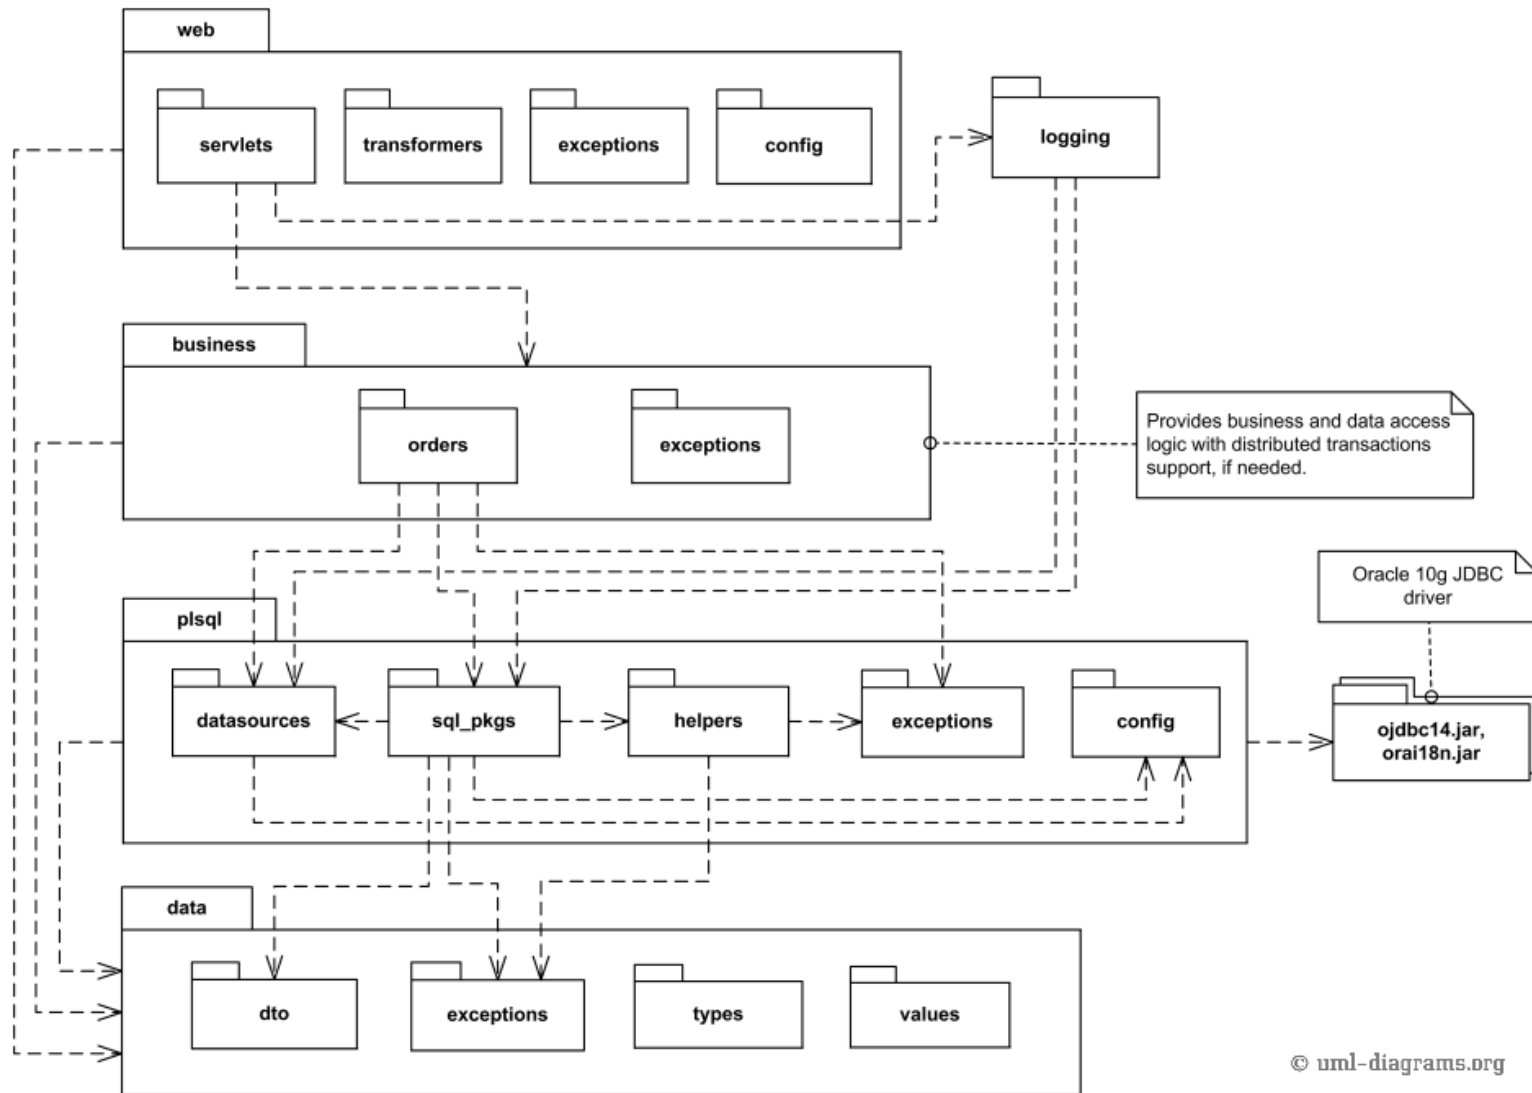


MULTI-LAYERED WEB ARCHITECTURE UML PACKAGE DIAGRAM EXAMPLE

Dependencies between packages are created in such a way as to avoid circular dependencies between packages.

Higher level packages depend on lower level packages.

Packages belonging to the same level could depend on each other. Data transfer objects and common exceptions are used by packages at higher levels.

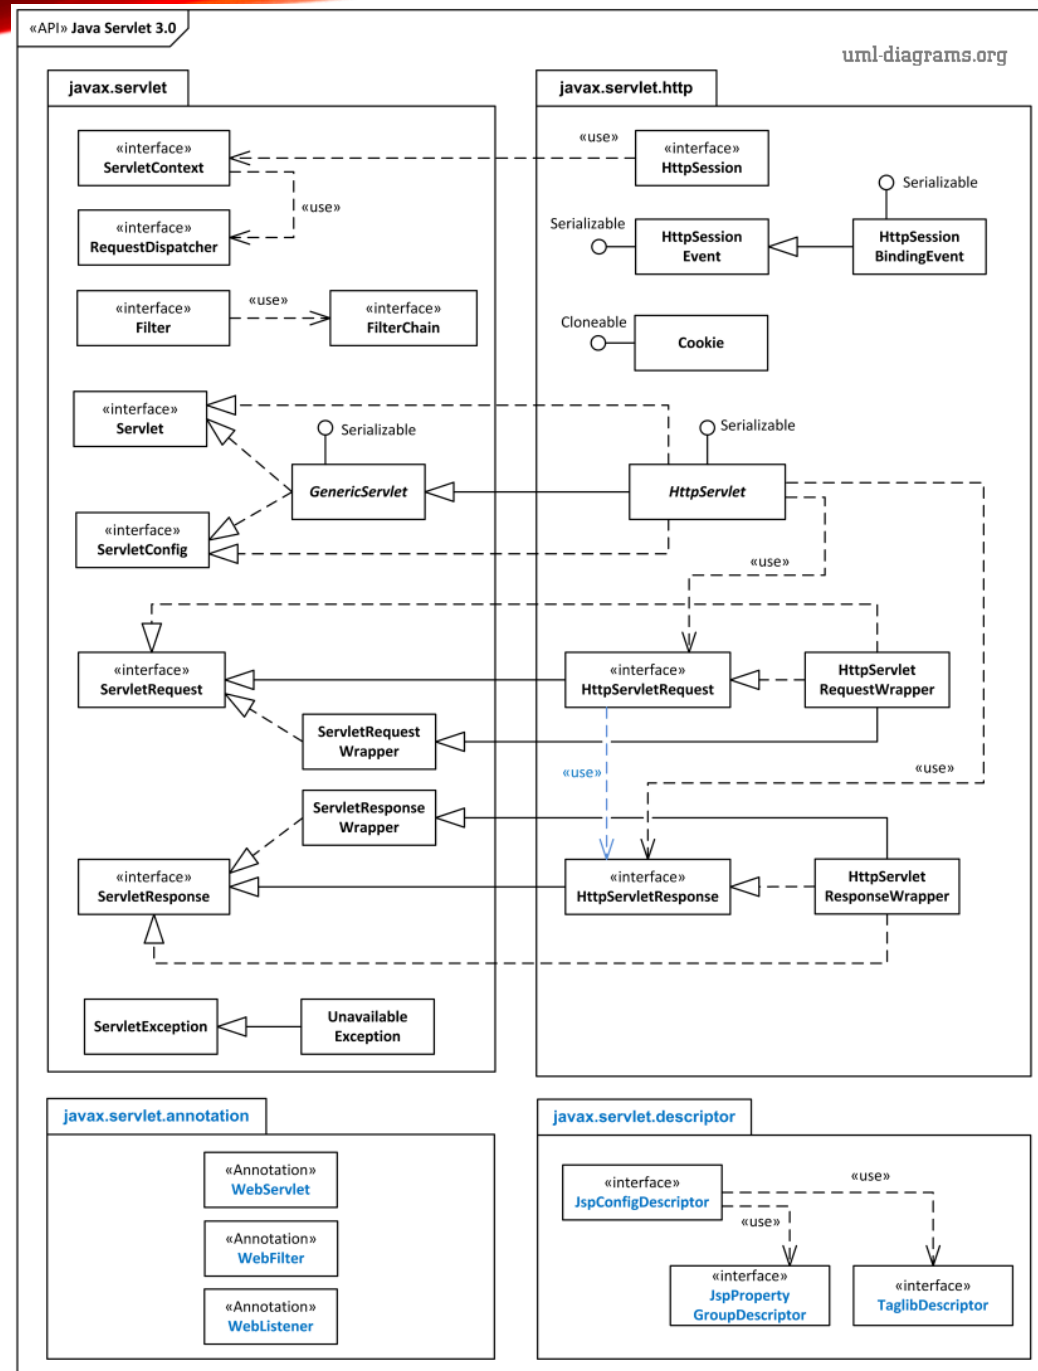


JAVA™ SERVLET 3.0 API

The first step can be Package Diagram
UML 2.5 doesn't have standard
stereotype to support modeling of APIs

Java Servlet 3.0 API consists of four
packages:

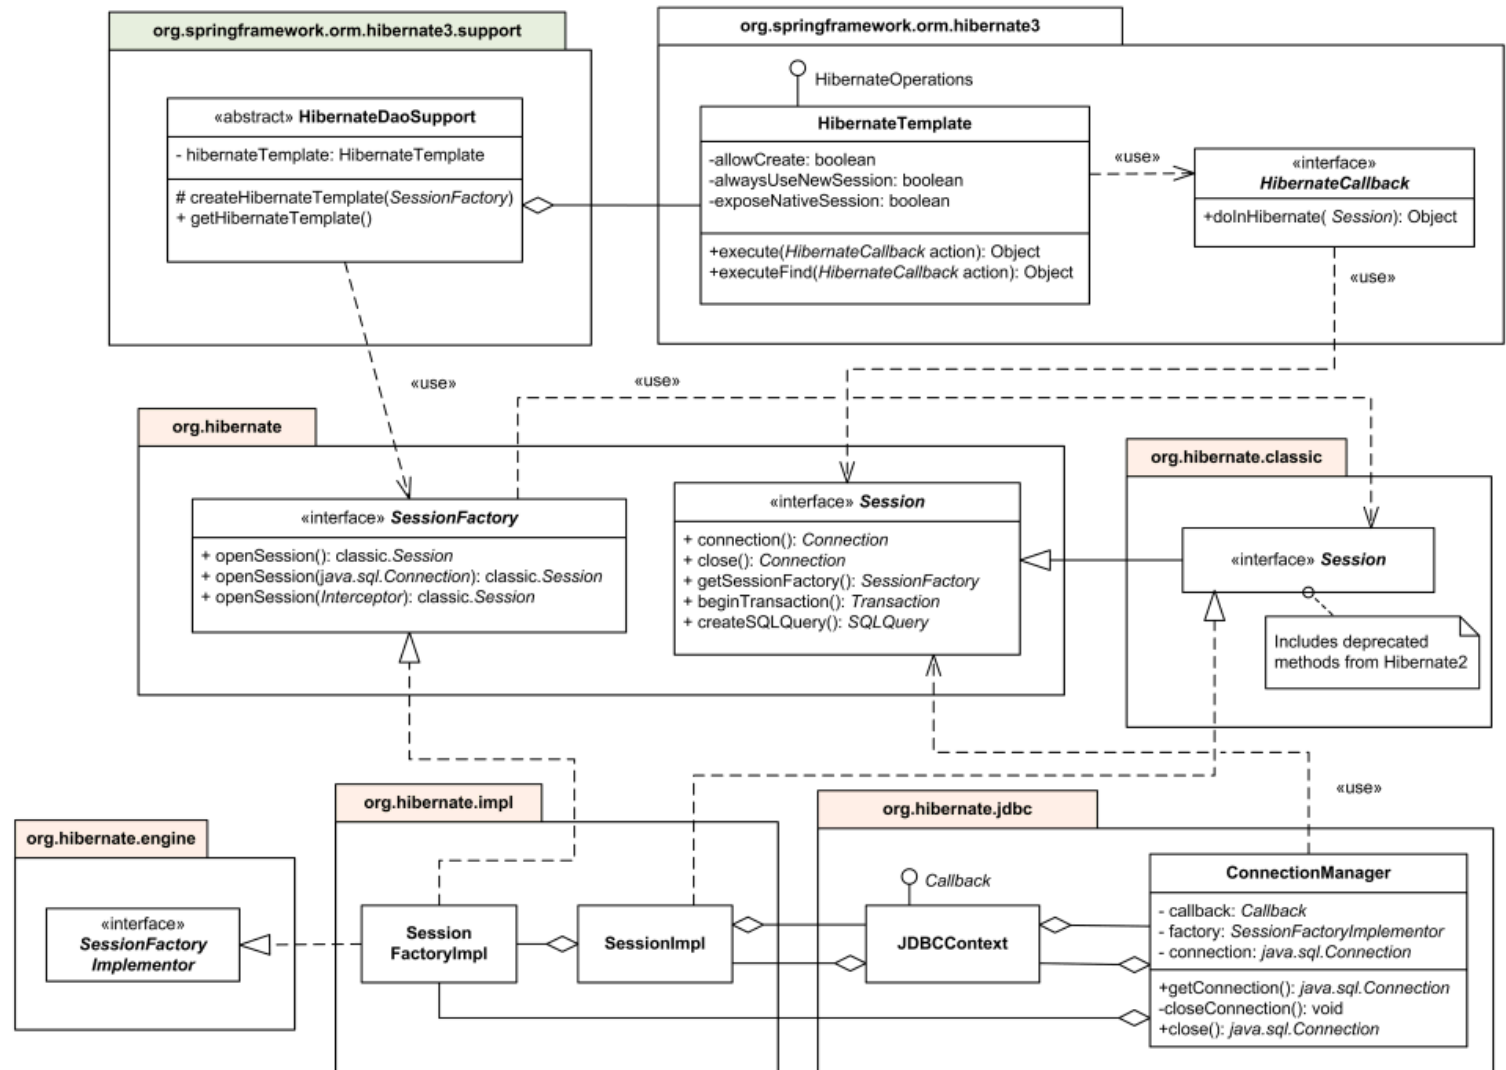
- **javax.servlet**,
- **javax.servlet.http**,
- **javax.servlet.annotation**, and
- **javax.servlet.descriptor**.



SPRING AND HIBERNATE CLASSES UML PACKAGE DIAGRAM EXAMPLE

Support packages for object relational mappers comply with Spring's generic transaction and DAO exception hierarchies.

Spring framework hibernate3 packages use several Hibernate packages.



PLANTUML

@startuml

left to right direction

skinparam packageStyle rectangle

actor customer

actor clerk

rectangle checkout {

customer -- (checkout)

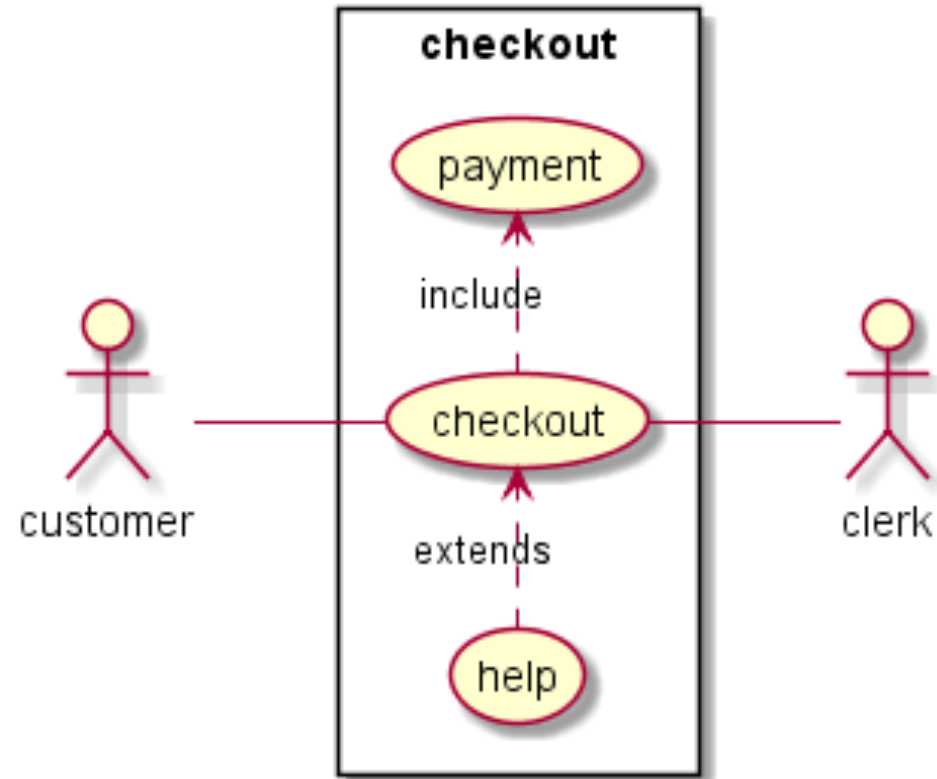
(checkout) .> (payment) : include

(help) .> (checkout) : extends

(checkout) -- clerk

}

@enduml




<https://plantuml.com/use-case-diagram>

PLATFORM AS A SERVICE



WHAT IS PAAS?

- PaaS is a framework for developers to create fully-customized applications with limited management responsibilities in the cloud.
- Platform and environment to allow developers to build applications and services
- Simply accessible via their web browser



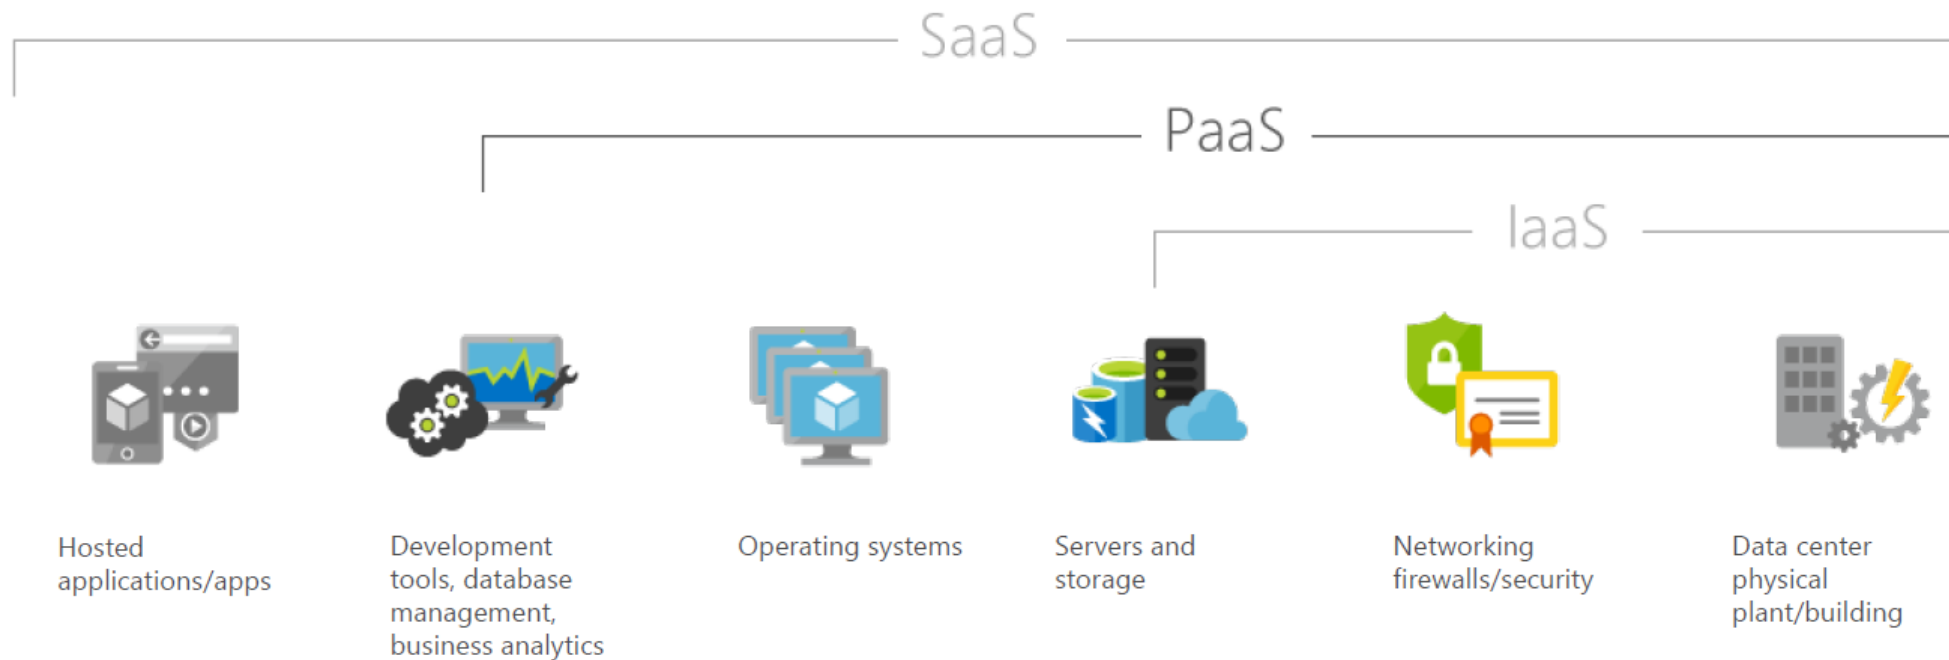
HOW DOES IT WORK?

- This is important for enterprise businesses seeking agility as they implement a DevOps approach to software development.
- The bottom line is that Paas requires fewer management responsibilities, allowing for greater focus to be placed on development.
- Creation of Software Applications
- Pay-per-use
- Choice of features
- Management and support
- Automatic upgrades

BENEFITS TO APPLICATION DEVELOPERS

- No investment in physical infrastructure
- Make development possible for 'non-experts'
- Flexibility & Adaptability
- Teams in various location can work together
- Security

PAAS IN AZURE



PAAS IN AZURE

- Web Apps
- Mobile Apps
- Logic Apps
- Functions
- Web Jobs

APPLICATION TIER

Monolithic vs Microservices




APPLICATION TIER

A monolithic application puts all its functionality into a single process...



... and scales by replicating the monolith on multiple servers



<http://martinfowler.com/articles/microservices.html>

**Problems with
Monolithic
approach?**

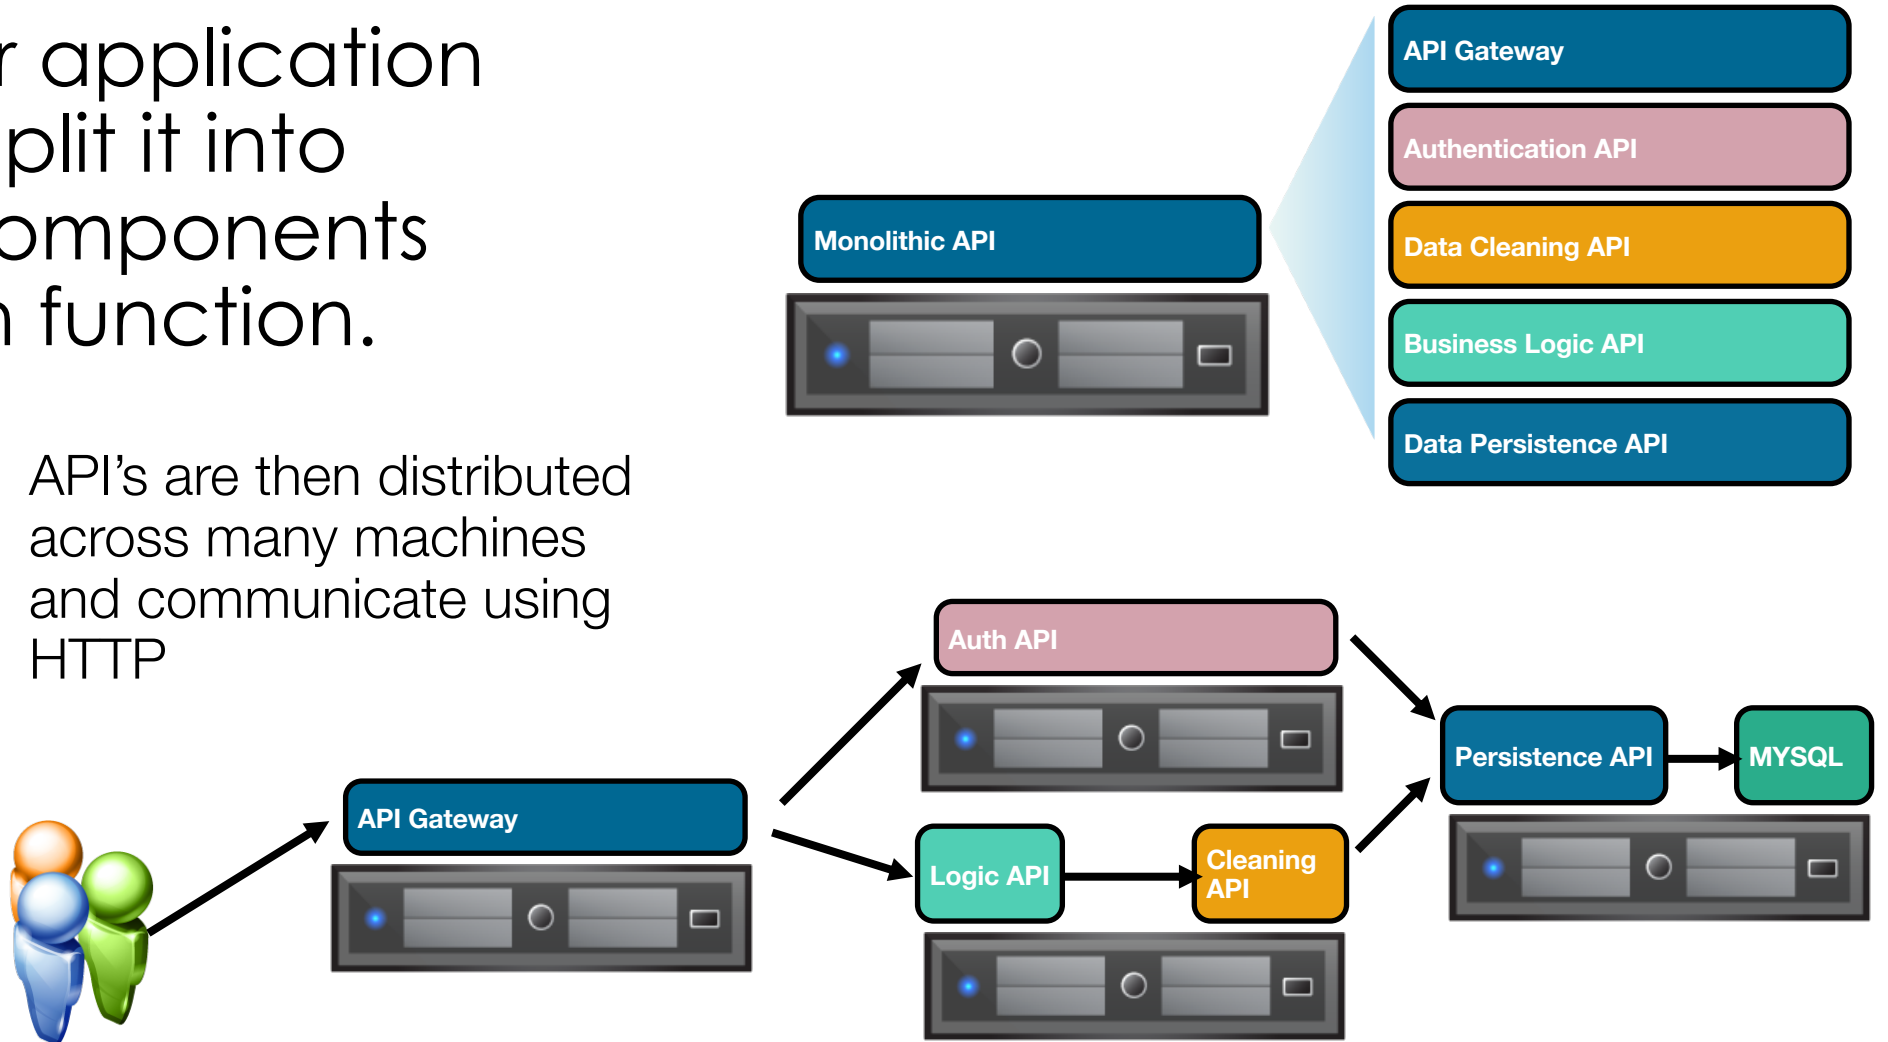
MONOLITHIC CHALLENGES

- **Scalability:** Need to possibly use both up and out to reach performance goals. Hard to scale things like databases.
- **Reliability:** A fault/memory leak in a single place can crash the entire app.
- **Orchestration:** You need to rebuild and deploy the entire application every time you make a change.
- **Code Complexity:** Code turns into spaghetti due to too many things happening at once. Hard to refactor features.
- **Upgradeability:** Moving to newer tech stacks requires converting the entire app at once.

MICROSERVICES

- Take your application API and split it into smaller components based on function.

API's are then distributed across many machines and communicate using HTTP




MICROSERVICES


A monolithic application puts all its functionality into a single process...




... and scales by replicating the monolith on multiple servers



A microservices architecture puts each element of functionality into a separate service...




... and scales by distributing these services across servers, replicating as needed.




Read more: <https://martinfowler.com/articles/microservices.html>

Challenges with Microservices approach?

*A microservices architecture puts
each element of functionality into a
separate service...*



*... and scales by distributing these services
across servers, replicating as needed.*



Read more: <https://martinfowler.com/articles/microservices.html>

MICROSERVICES CHALLENGES

- **Discovery:** how to find a service you want?
- **Scalability:** how to replicate services for speed?
- **Openness:** how to agree on a message protocol?
- **Fault tolerance:** how to handle failed services?

All distributed systems face these challenges, microservices just increases the scale and diversity...

NETFLIX

- 20th most popular website according to Alexa
- Zero of their own servers
 - All infrastructure is on AWS (2016-2018)
 - Recently starting to build out their own Content Delivery Network

NETFLIX is **15%**
of the total downstream volume of traffic
across the entire internet

NETFLIX

- One of the first to really push microservices
 - Known for their DevOps
 - Fast paced, frequent updates, must always be available
- 700+ microservices
- Deployed across 10,000s of VMs and containers

Netflix ecosystem

100s of microservices

1000s of daily production changes

10,000s of instances

100,000s of customer interactions per minute

1,000,000s of customers

1,000,000,000s of metrics


10,000,000,000 hours of streamed

10s of operations engineers

Netflix tech talk: <https://www.youtube.com/watch?v=CZ3wluvmHeM>

NETFLIX “DEATHSTAR”

- Microservice architecture results in a extremely distributed application
 - Can be very difficult to manage and understand how it is working at scale
- What if there are failures?
- How to know if everything is working correctly?



NETFLIX CHAOS MONKEY

- Idea: If my system can handle failures, then I don't need to know exactly how all the pieces themselves interact!
- Chaos Monkey:
 - Randomly terminate VMs and containers in the production environment
 - Ensure that the overall system keeps operating
 - Run this 24/7

<http://principlesofchaos.org/>



Make failures the common case, not an unknown!

MICROSERVICE ARCHITECTURE MODELING WITH DOMAIN DRIVEN DESIGN



Domain Events:

A domain event

- is something of interest that has happened to an aggregate
- should be expressed in past tense
- typically represents state change
- should be represented by a class in the domain model
- may be registered in an event class hierarchy

<https://www.umlzone.com/videos/modeling-microservices-with-domain-driven-design/>

SERVERLESS COMPUTING



SERVERLESS COMPUTING


- Trendy architecture that improves the agility of microservices
- What does “serverless” mean?



AWS Lambda



APACHE
OpenWhisk™




Azure
Functions



Google Cloud Functions


SERVERLESS COMPUTING (FAAS)

- Trendy architecture that improves the agility of microservices
- What does “serverless” mean?
- You still need a server!
- BUT, your services will not always be running
- Key idea: only **instantiate** a service when a user makes a request for that functionality; **charge** the cloud customer based on number of requests
- How will this work for stateful vs stateless services?

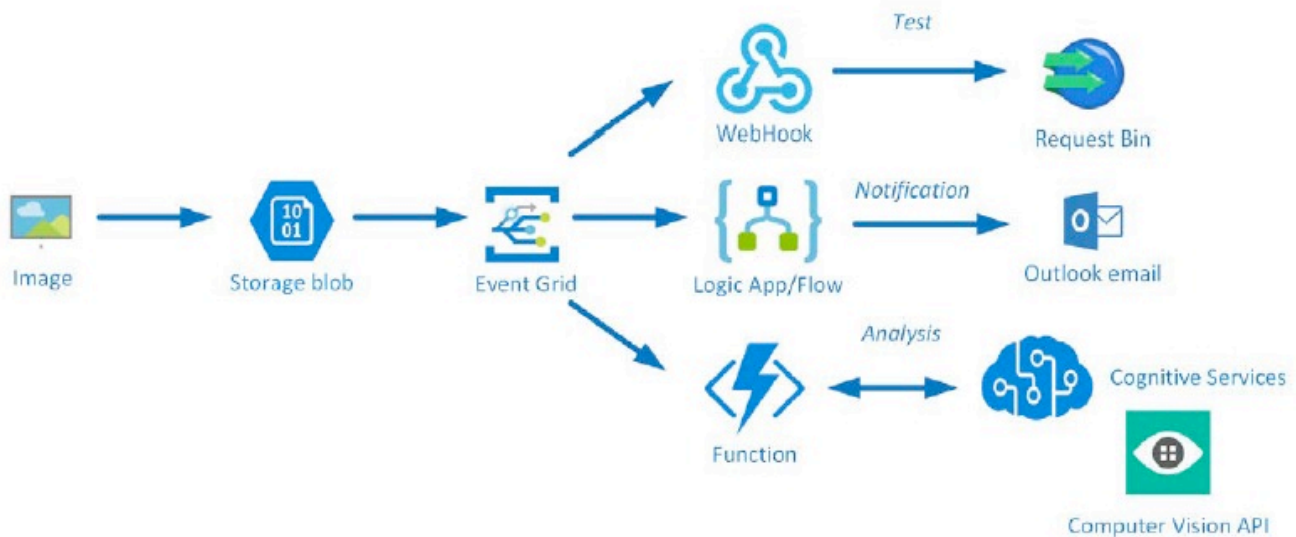


USAGE EXAMPLES OF SERVERLESS

- IoT (Weather stations)




- Data processing (Image Manipulation)



USAGE EXAMPLES OF SERVERLESS

- Web Application (Weather stations)



- Chatbots (Event-Driven Architecture for chatbot)



SERVERLESS IN AZURE



Storage



Functions



Logic App



Media Services



Cortana
Management Suite



App Service



Document DB



CDN



Machine Learning



OMS
Management Suit



Visual Studio
Services



Traffic
Manager



Active Directory



Key Vault



App Insights



Cognitive Services



Embedded Power BI



Search



IoT Hub



Service Bus



Notification Hub



Stream Analytics



Hockey App



Scheduler

SERVERLESS STARTUP

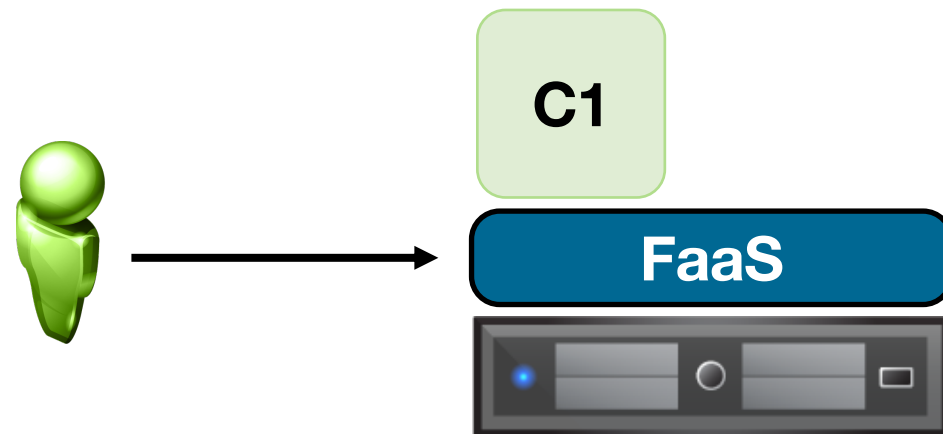
- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed

No workload means no resources being used!



SERVERLESS STARTUP


- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed



Request arrives, start green container

SERVERLESS STARTUP

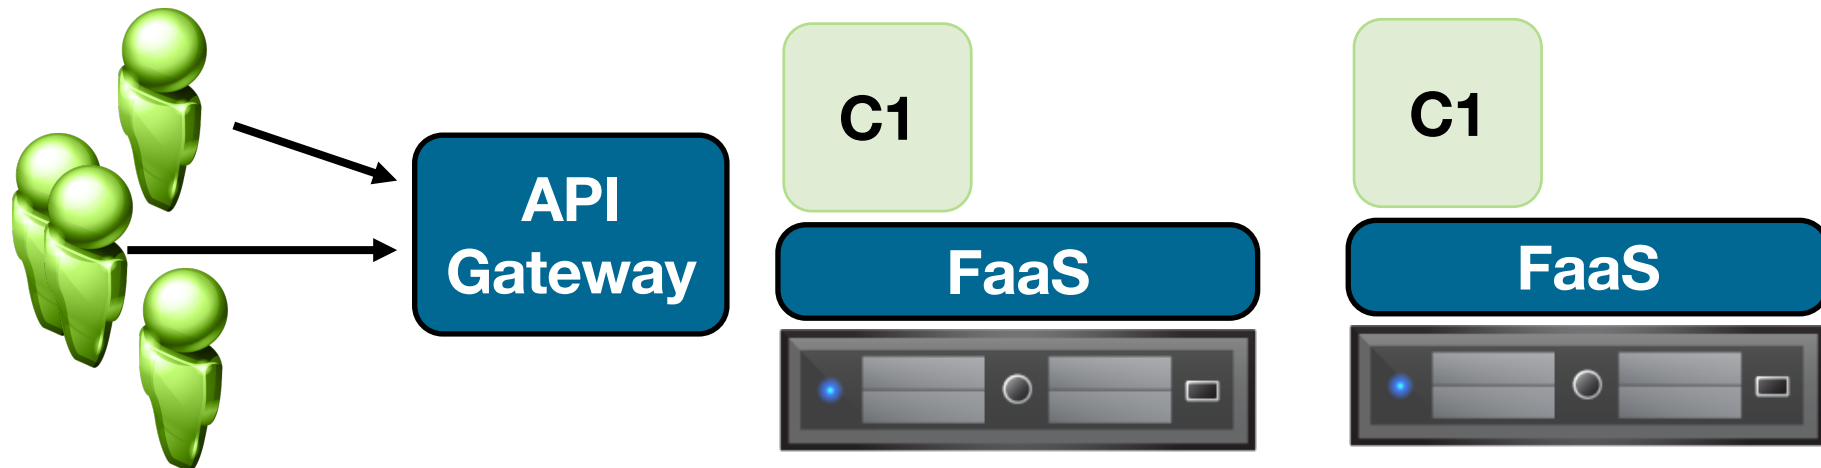
- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed



Reuse that container for subsequent requests

SERVERLESS STARTUP

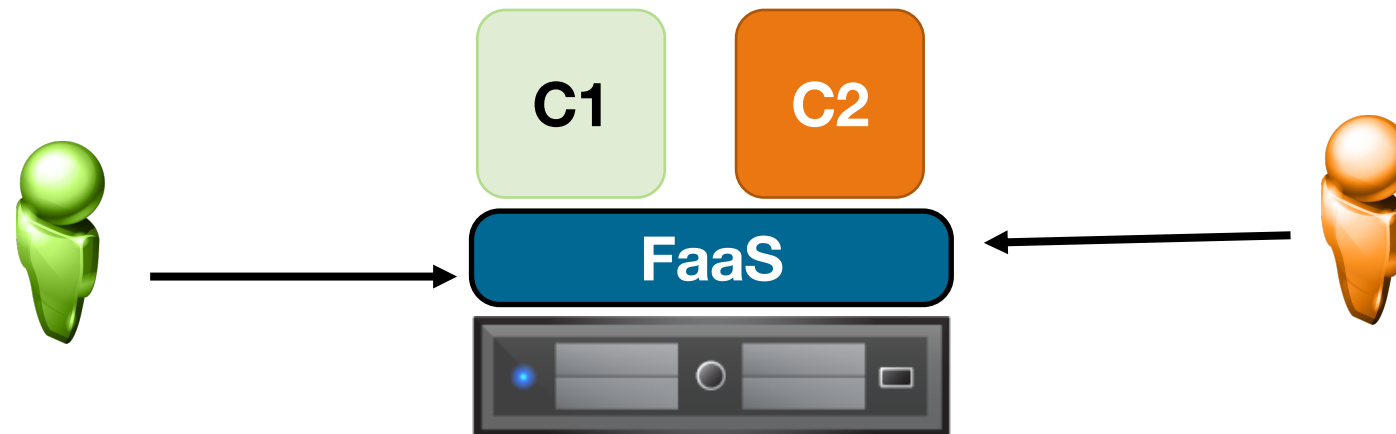
- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed



Add more replicas if workload exceeds capacity

SERVERLESS STARTUP

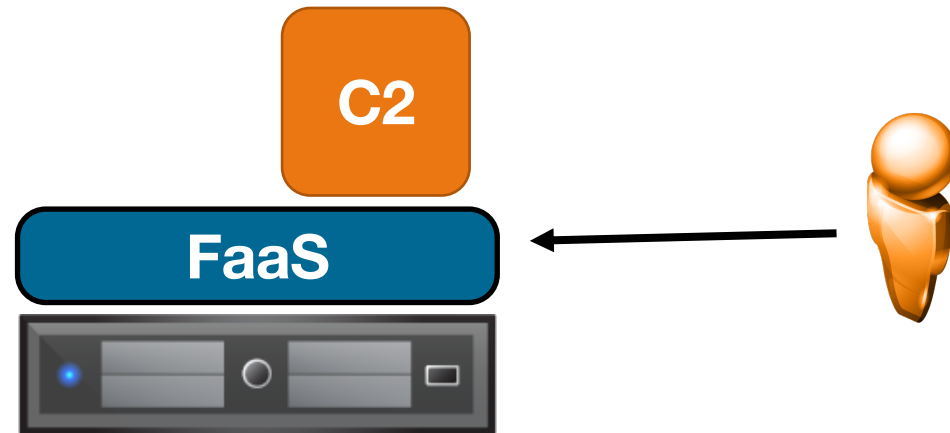
- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed



Start new container if user needs a different function

SERVERLESS STARTUP

- Function as a Service platforms (Lambda, GCF, Azure Functions, etc)
 - Define a stateless “function” to execute for each request
 - A container will be instantiated to handle the first request
 - The same container will be used until it times out or is killed

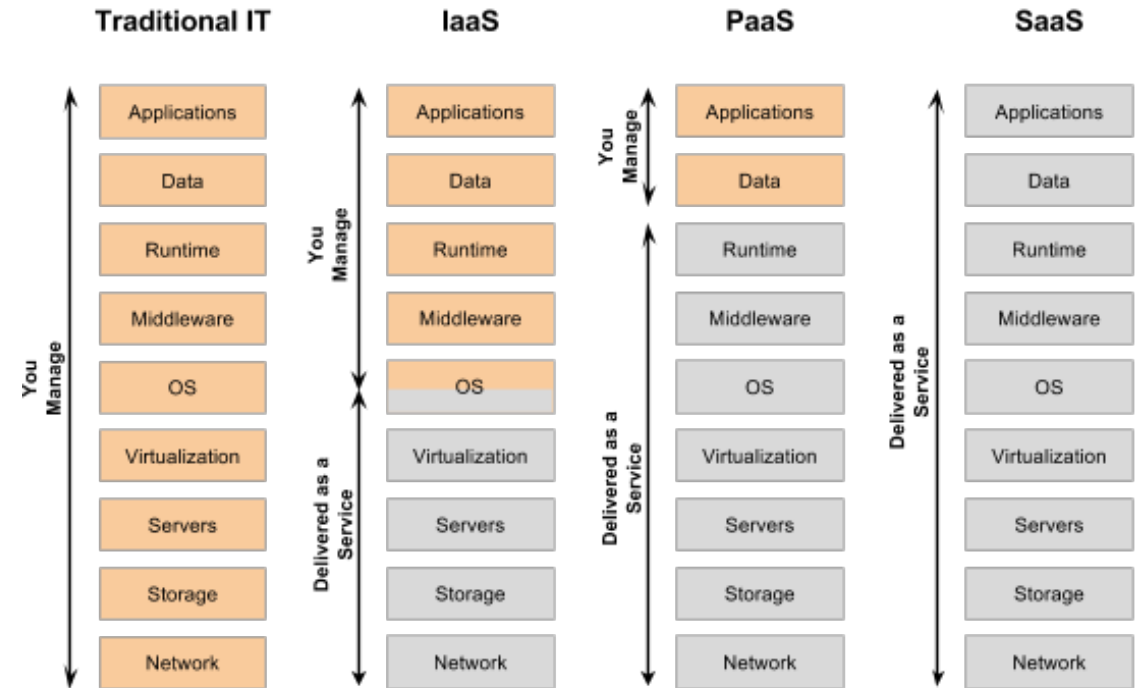
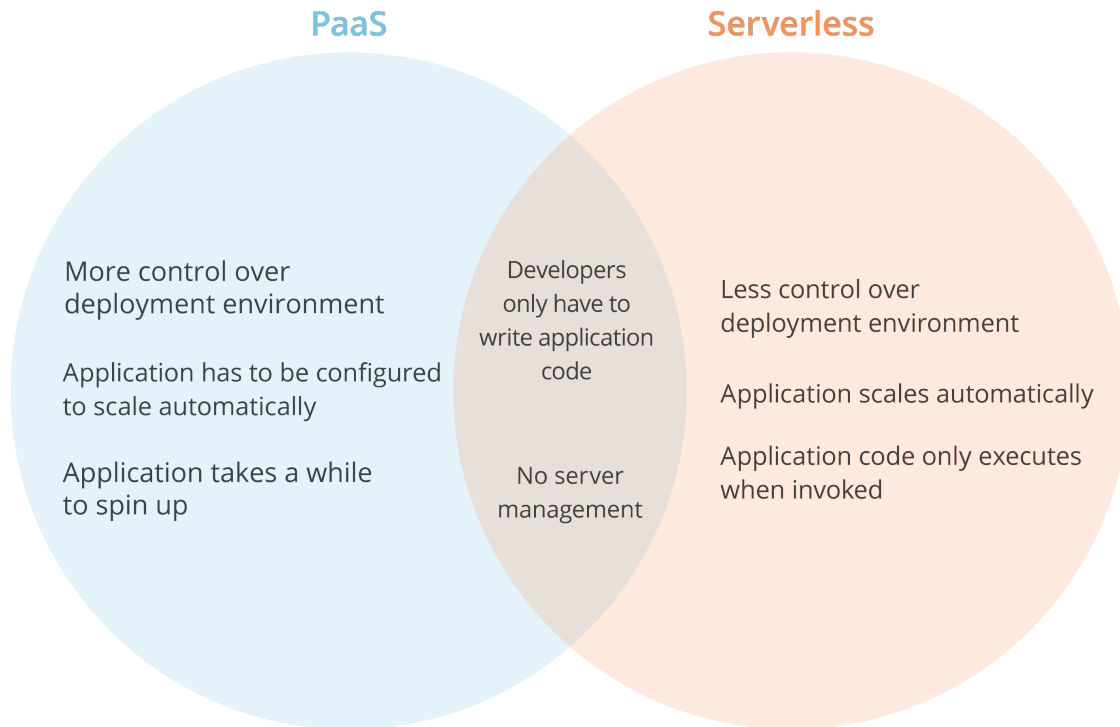


Stop old containers once not in use

SERVERLESS PROS/CONS

- Benefits:
 - Simple for developer when auto scaling up
 - Pay for exactly what we use (at second granularity)
 - Efficient use of resources (auto scale up and down based on requests)
 - don't worry about reliability/server management at all
- Drawbacks:
 - Limited functionality (stateless, limited programming model)
 - High latency for first request to each container
 - Some container layer overheads plus the lambda gateway and routing overheads
 - Potentially higher and unpredictable costs
 - Difficult to debug / monitor behavior
 - Security

SERVERLESS VS PaaS?



SERVERLESS DEMO

