

# TCSS 562: SOFTWARE ENGINEERING FOR CLOUD COMPUTING

## Cloud Computing: Fundamental Concepts and Models

Wes J. Lloyd  
School of Engineering and Technology  
University of Washington - Tacoma



The illustration on the right side of the slide depicts cloud computing concepts. The top part shows a blue cloud with lines connecting it to a laptop, a smartphone, and a tablet. The bottom part shows a server rack with multiple monitors displaying data.

1

## FEEDBACK FROM 10/14

- Perspective on material: 6.667 (→ *mostly new to me*)
- Pace: 5.333 (~ just right)
- 18 respondents
- What is a billing model?
- In tutorial 3, along with CSV output, we need to upload even the graphs. What exactly are the graphs we should upload/attach?
  - See bottom of page 10 for explanation.
  - Looking for HTML output pasted into DOC/PDF ideally

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.2
------------------	--	------

2


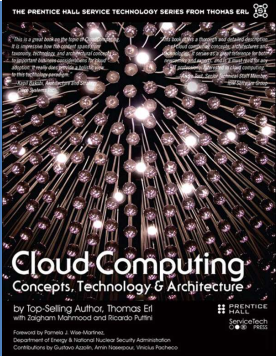
FEEDBACK - 2

- **What is the format of questions on the midterm exam? Are questions objective or subjective?**
  - A practice midterm will be given as an in class activity prior to the midterm to practice question format
  - There are objective questions
  - There are also questions that ask about trade-offs of alternatives
    - i.e. weigh and compare differences
- **What is the time duration of the midterm?**
  - Full 2 hours is permitted, .....
- **What does m-bound and d-bound mean?**
  - M-bound: performance bottleneck is the soil erosion model
  - D-bound: performance bottleneck is the relational database (pgsql)

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.3
------------------	--	------

3

CHAPTER 4: FUNDAMENTAL CONCEPTS AND MODELS



October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.4
------------------	--	------

4

OBJECTIVES

- From: Cloud Computing Concepts, Technology & Architecture:
- Cloud Computing Concepts and Models
  - Roles and boundaries
  - Cloud characteristics
  - **Cloud delivery models**
  - Cloud deployment models

October 16, 2019





TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.5

5

CLOUD DELIVERY MODELS

- What is the appropriate level of abstraction?
- How should applications be deployed?
  - IaaS, PaaS, SaaS, DbaaS, FaaS
- How do we ensure Quality-of-Service?
  - Performance, Availability, Responsiveness, Fault Tolerance
- How is scalability provided?
- How do we minimize hosting costs?
  - How do we estimate hosting costs?

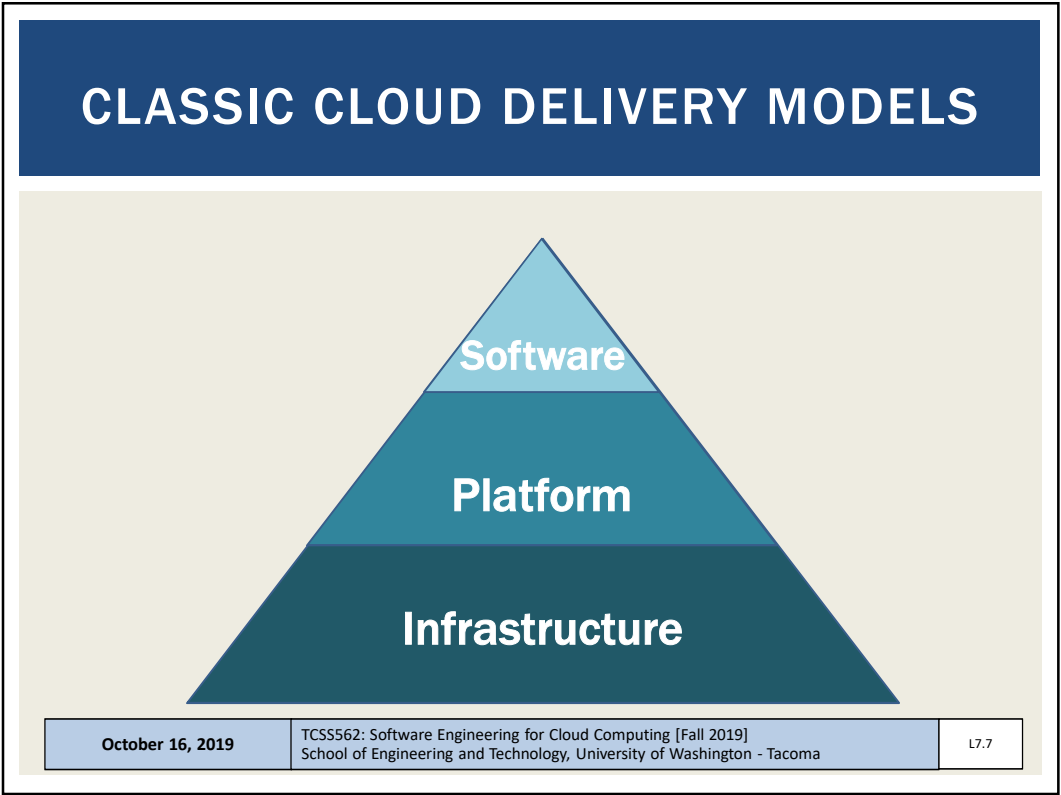


October 16, 2019

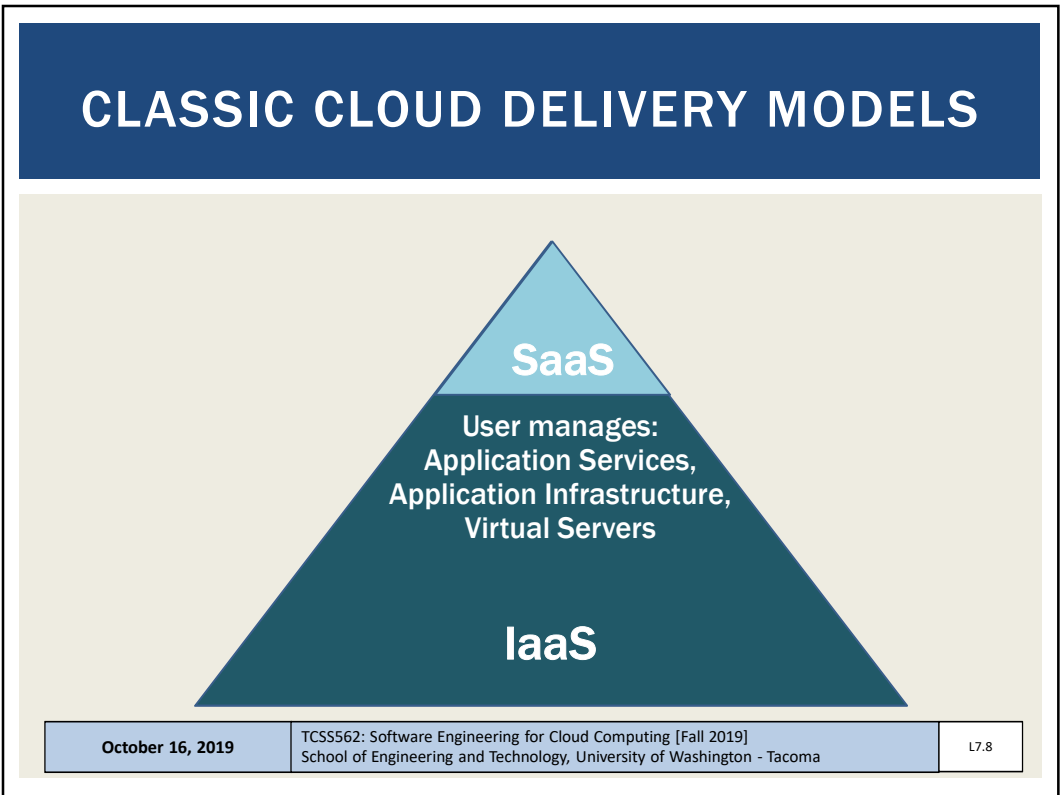
TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.6

6



7



8

# CLASSIC CLOUD DELIVERY MODELS

A pyramid diagram illustrating the hierarchy of cloud delivery models. The pyramid is divided into three horizontal layers. The top layer is light blue and labeled 'SaaS'. The middle layer is a medium blue and labeled 'PaaS', with the text 'User manages: Application Services' positioned just above it. The bottom layer is a dark blue and labeled 'IaaS'. The entire pyramid is set against a light beige background.

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.9
------------------	--	------

9


# CLASSIC CLOUD DELIVERY MODELS

A pyramid diagram illustrating the hierarchy of cloud delivery models. The pyramid is divided into three horizontal layers. The top layer is light blue and labeled 'SaaS'. The middle layer is a medium blue and labeled 'PaaS'. The bottom layer is a dark blue and labeled 'IaaS'. The entire pyramid is set against a light beige background.

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.10
------------------	--	-------

10


EXAMPLE CLOUD SERVICES



SAAS  
Software as a Service

Email  
CRM  
Collaborative  
ERP


CONSUME



PAAS  
Platform as a Service

Application Development  
Decision Support  
Web  
Streaming

BUILD ON IT



IAAS  
Infrastructure as a Service

Caching  
Legacy  
Networking  
Security  
File  
Technical  
System Mgmt

MIGRATE TO IT

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.11

11

END USER APPLICATIONS

Many different "cloud" providers

Software-as-a-Service

Finance & Accounting

Content Management

Vertical

Enterprise Social Media

Marketing Analytics

Retail & E-Commerce

Collaboration

Business Intelligence

Ad Tech

Many cloud providers are also cloud consumers

Cloud Foundry

Infrastructure-as-a-Service

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.12

12

Slides by Wes J. Lloyd

L7.6

INFRASTRUCTURE-AS-A-SERVICE

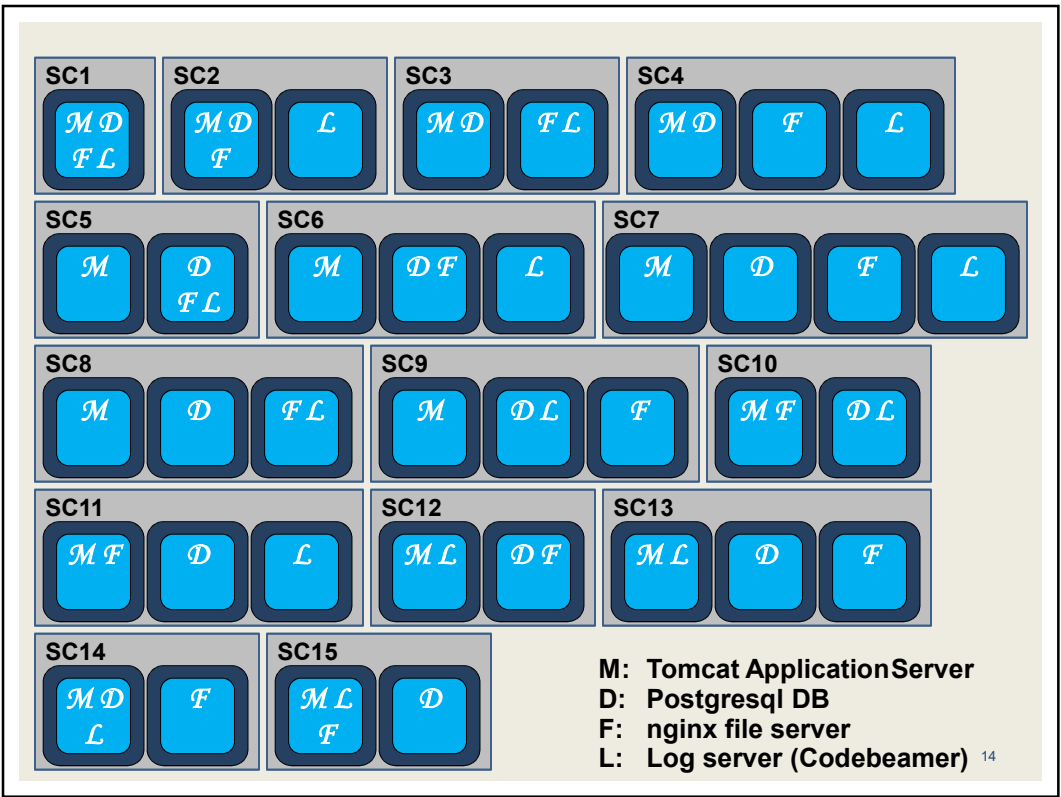
- Compute resources, on demand, as-a-service
  - Generally raw “IT” resources
  - Hardware, network, containers, operating systems
- Typically provided through virtualization
- Generally not-preconfigured
- Administrative burden is owned by cloud consumer
- Best when high-level control over environment is needed
- Scaling is generally **not** automatic...
- Resources can be managed in bundles
- AWS CloudFormation: Allows specification in JSON/YAML of cloud infrastructures

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.13

13



14

SC1

M D

F L

SC2

M D

F

L

SC3

M D

F L

SC4

M D

F

L

Bell's Number:

k: number of ways  
n components can be  
distributed across containers

n	k
4	15
5	52
6	203
7	877
8	4,140
9	21,147
n	...

SC14

M D

L

F

SC15

M L

F

D

M: Tomcat ApplicationServer

D: Postgresql DB

F: nginx file server

L: Log server (Codebeamer)

15

SC1

M D

F L

SC2

M D

F

L

SC3

M D

F L

SC4

M D

F

L

Component Composition Example

An application with 4 components has 15 compositions

One or more component(s) deployed to each VM

Each VM launched to separate physical machine

SC5

M

D

SC6

M

D F

L

SC7

M

D

F

L

SC14

M D

L

F

SC15

M L

F

D

M: Tomcat ApplicationServer


D: Postgresql DB

F: nginx file server


L: Log server (Codebeamer)

16







17




18



19



20



21

# PLATFORM-AS-A-SERVICE

- Predefined, ready-to-use, hosting environment
- Infrastructure is further obscured from end user
- Scaling and load balancing may be automatically provided and automatic
- Variable to no ability to influence responsiveness

■ Examples:

- Google App Engine
- Heroku
- AWS Elastic Beanstalk
- AWS Lambda (FaaS)

```
graph LR; Clients --> LB[Load Balancer]; LB --> Servers;
```

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.22

22


USES FOR PAAS

- Cloud consumer
  - Wants to extend on-premise environments into the cloud for “web app” hosting
  - Wants to entirely substitute an on-premise hosting environment
  - Cloud consumer wants to become a cloud provider and deploy its own cloud services to external users
- PaaS spares IT administrative burden compared to IaaS

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.23
------------------	--	-------


23

SERVERLESS COMPUTING



**What is serverless?**

Build and run applications without thinking about servers



October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.24
------------------	--	-------

24

# SERVERLESS COMPUTING - 2

## Evolving to serverless

Physical servers in datacenters

Virtual servers in datacenters

Virtual servers in the cloud

## SERVERLESS

amazon web services

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.25

25

# SERVERLESS COMPUTING

Pay only for CPU/memory utilization

High Availability

Fault Tolerance

Infrastructure Elasticity

No Setup

Function-as-a-Service (FAAS)

26

## SERVERLESS COMPUTING

### Why Serverless Computing?

**Many features of distributed systems,  
that are challenging to deliver, are  
provided automatically**

*...they are built into the platform*

27

## SERVERLESS VS. FAAS


- **Serverless Computing**
- Refers to the avoidance of managing servers
- Can pertain to a number of “as-a-service” cloud offerings
- **Function-as-a-Service (FaaS)**
  - Developers write small code snippets (microservices) which are deployed separately
- **Database-as-a-Service (DBaaS)**
- **Container-as-a-Service (CaaS)**
- **Others...**
  
- **Serverless is a buzzword**
- **This space is evolving...**

October 16, 2019


TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.28

28




29



30

## OPEN SOURCE FAAS FRAMEWORKS



Nuclio graphic

- Deployable to Docker container(s) or a Kubernetes cluster
- Fission: <https://fission.io/>
- Kubeless: <https://kubeless.io/>
- Nuclio: <https://nuclio.io/>
- OpenFaaS: <https://www.openfaas.com/>


- Supports cloud native development principles
- Building a cloud application by adopting a “deploy it yourself” framework avoids vendor lock-in
- Requires common medium of Kubernetes


October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.31
------------------	--	-------

31

## AWS LAMBDA

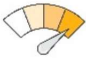
### Using AWS Lambda






#### Bring your own code

- Node.js, Java, Python, C#
- Bring your own libraries (even native ones)




#### Simple resource model

- Select power rating from 128 MB to 3 GB
- CPU and network allocated proportionately



#### Flexible use

- Synchronous or asynchronous
- Integrated with other AWS services



#### Flexible authorization

- Securely grant access to resources and VPCs
- Fine-grained control for invoking your functions

Images credit: aws.amazon.com

32



## FAAS PLATFORMS - 2

- New cloud platform for hosting application code
- Every cloud vendor provides their own:
  - AWS Lambda, Azure Functions, Google Cloud Functions, IBM OpenWhisk
- Similar to platform-as-a-service
- Replace opensource web container (e.g. Apache Tomcat) with abstracted vendor-provided **black-box** environment

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.33

33

## FAAS PLATFORMS - 3

- Many challenging features of distributed systems are provided automatically
- **Built into the platform:**
- Highly availability (24/7)
- Scalability
- Fault tolerance

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.34

34

# CLOUD NATIVE SOFTWARE ARCHITECTURE

- Every service with a different pricing model

Example: Weather Application

S3      API GATEWAY      LAMBDA      DYNAMODB

Front-end code for weather app hosted in S3      User clicks on link to get local weather information      App makes REST API call to endpoint      Lambda is triggered      35° C      Lambda runs code to retrieve local weather information and returns data back to user

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.35
------------------	--	-------

35

# IAAS BILLING MODELS

- Virtual machines as-a-service at ¢ per hour
- No premium to scale:

1000 computers

@

1 hour

=

1 computer

@

1000 hours
- Illusion of infinite scalability to cloud user
- As many computers as you can afford
- Billing models are becoming increasingly granular
  - By the minute, second, 1/10th sec
- Auction-based instances:  
Spot instances →

Spot Instance Pricing History

Product: Linux/UNIX (Amazon VPC)      Instance Type: c4.xlarge

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.36
------------------	--	-------

36

IAAS VS. FAAS COMPUTING  
BILLING MODELS

- AWS Lambda Pricing
- FREE TIER:
  - first 1,000,000 function calls/month → FREE
  - first 400,000 GB-sec/month → FREE
- Afterwards: *obfuscated pricing (AWS Lambda):*
  - \$0.0000002 per request
  - \$0.000000208 to rent 128MB / 100-ms

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.37

37

WEBSERVICE HOSTING EXAMPLE

- Workload: 1-month continuous 1-second service calls that fully utilize 3GB of RAM and two CPU cores
- ON AWS Lambda
  - Each service call: 100% of 1 CPU-core  
100% of 3GB of memory
  - Workload: 2 continuous client threads
  - Duration: 1 month (30 days)
- ON AWS EC2:
  - Amazon EC2 c4.large 2-vCPU VM@3.75GB
  - Hosting cost: \$72/month  
c4.large: 10¢/hour, 24 hrs/day x 30 days
- **How much would hosting this workload cost on AWS Lambda?**

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.38

38

PRICING OBFUSCATION

■ Workload:7,776,000 GB-sec

■ FREE:-400,000 GB-sec

■ Ch...

■ M...

■ In AWS EC2:\$72.00

■ FF AWS Lambda:\$123.28

■ Charge...

■ Calls:\$0.32

■ Total:\$123.28

■ BREAK-EVEN POINT = ~4,319,136 GB-sec-month

Worst-case scenario = ~1.7x

For compute only, not considering cost of function calls = ~16.7 days

39

FAAS PRICING

■ Break-even point is the point where renting VMs or deploying to a serverless platform (e.g. Lambda) is exactly the same.

■ Our example is for one month

■ Could also consider one day, one hour, one minute

- What factors influence the break-even point for an application running on AWS Lambda?
- What scenario would result in a 1-day break-even point where pricing for IaaS=FaaS?

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.40

40

FACTORS IMPACTING PERFORMANCE OF  
FAAS COMPUTING PLATFORMS

- Infrastructure elasticity
- Load balancing
- Provisioning variation
- Infrastructure retention: COLD vs. WARM
  - Infrastructure freeze/thaw cycle
- Memory reservation
- Service composition

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.41

41

FAAS CHALLENGES

- Outline:
- Vendor architectural lock-in – how to migrate?
- Pricing obfuscation – is it cost effective?
- Memory reservation – how much to reserve?
- Service composition – how to compose software?
- Infrastructure freeze/thaw cycle – how to avoid?

October 16, 2019

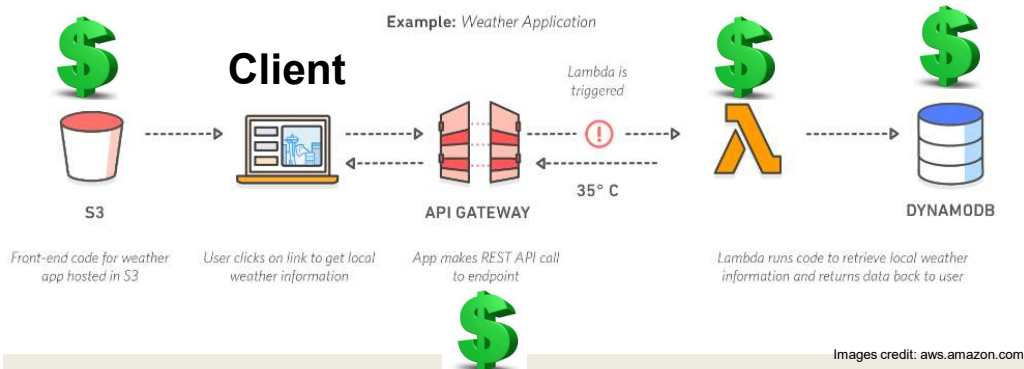
TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.42

42

## VENDOR ARCHITECTURAL LOCK-IN

- Cloud native (FaaS) software architecture requires external services/components



- Increased dependencies → increased hosting costs

43

## PRICING OBFUSCATION

- VM pricing:** hourly rental pricing, billed to nearest second is intuitive...

- FaaS pricing:**

### AWS Lambda Pricing

**FREE TIER:** first 1,000,000 function calls/month → FREE  
first 400 GB-sec/month → FREE

- Afterwards:** \$0.0000002 per request  
\$0.000000208 to rent 128MB / 100-ms


October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.44

44

## MEMORY RESERVATION QUEST



- Lambda memory reserved for functions
- UI provides “slider bar” to set function’s memory allocation
- Resource capacity (CPU, disk, network) coupled to slider bar:  
“every *doubling* of memory, *doubles* CPU...”
- But how much memory do model services require?

**Basic settings**

Memory (MB) Info  
Your function is allocated CPU proportional to the memory configured.

1536 MB

Timeout Info  
3 min 0 sec

Description

**Performance**


October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.45
------------------	--	-------

45

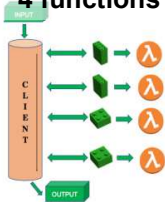
## SERVICE COMPOSITION

- How should application code be composed for deployment to serverless computing platforms?

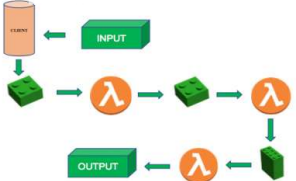
**Monolithic Deployment**



**Client flow control, 4 functions**




**Server flow control, 3 functions**



- Recommended practice:  
Decompose into many microservices
- Platform limits: code + libraries ~250MB
- How does composition impact the number of function invocations, and memory utilization?

**Performance**



46

## INFRASTRUCTURE FREEZE/THAW CYCLE

- Unused infrastructure is deprecated
  - *But after how long?*
- Infrastructure: VMs, “containers”
- Provider-COLD / VM-COLD
  - “Container” images - built/transferred to VMs
- Container-COLD
  - Image cached on VM
- Container-WARM
  - “Container” running on VM

**Performance**






Image from: Denver7 – The Denver Channel News

47



## FUNCTION-AS-A-SERVICE

AWS  
Lambda  
Demo

48

48



## SOFTWARE-AS-A-SERVICE

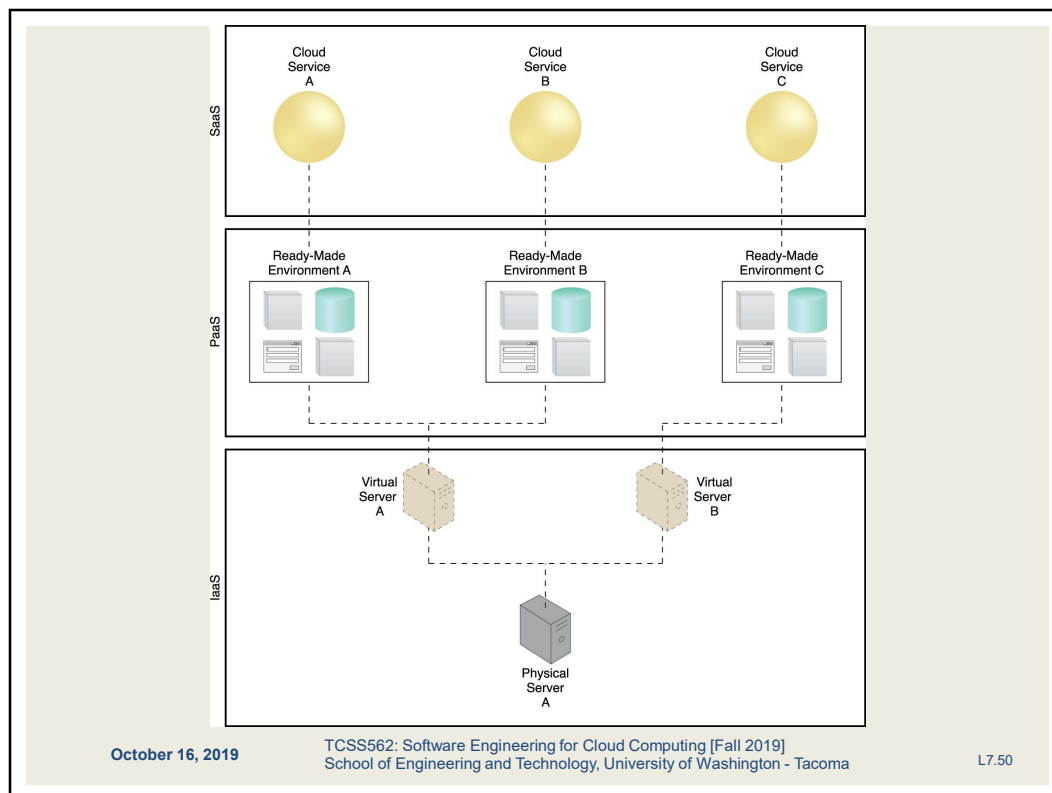
- Software applications as shared cloud service
- Nearly all server infrastructure management is abstracted away from the user
- Software is generally configurable
- SaaS can be a complete GUI/UI based environment
- Or UI-free (database-as-a-service)
- SaaS offerings
  - Google Docs
  - Office 365
  - Cloud9 Integrated Development Environment
  - Salesforce

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.49

49



50

CONTAINER-AS-A-SERVICE

- Cloud service model for deploying application containers (e.g. Docker) to the cloud
- Deploy containers without worrying about managing infrastructure:
  - Servers
  - Or container orchestration platforms
  - Container platform examples: Kubernetes, Docker swarm, Apache Mesos/Marathon, Amazon Elastic Container Service
  - Container platforms support creation of container clusters on the using cloud hosted VMs
- CaaS Examples:
  - AWS Fargate
  - Azure Container Instances
  - Google Cloud Run
  - Open Source – deploy on your datacenter: Knative (led by Google)

October 16, 2019

TCCS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.51

51

OTHER CLOUD SERVICE MODELS

- IaaS
  - Storage-as-a-Service
- PaaS
  - Integration-as-a-Service
- SaaS
  - Database-as-a-Service
  - Testing-as-a-Service
  - Model-as-a-Service
- ?
  - Security-as-a-Service
  - Integration-as-a-Service

October 16, 2019

TCCS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L10.52

52

OBJECTIVES

- Cloud Computing Concepts and Models
  - Roles and boundaries
  - Cloud characteristics
  - Cloud delivery models
  - Cloud deployment models

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.53

53

CLOUD DEPLOYMENT MODELS

- Distinguished by ownership, size, access
- Four common models
  - Public cloud
  - Community cloud
  - Hybrid cloud
  - Private cloud

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.54

54

# PUBLIC CLOUDS

The diagram illustrates the public cloud model. At the top, a collection of cloud icons represents different service providers: Salesforce, Microsoft, Google, Yahoo, Amazon, Zoho, and Rackspace. Below these, three server racks are shown, representing the organizations that use these services. Arrows point from the organizations up to the cloud providers, indicating the flow of data and services.

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.55
------------------	--	-------

55

# COMMUNITY CLOUD

- Specialized cloud built and shared by a particular community
- Leverage economies of scale within a community
- Research oriented clouds
- Examples:
  - Bionimbus - bioinformatics
  - Chameleon
  - CloudLab

The diagram illustrates the community cloud model. A single cloud icon is shown, containing several server racks and storage cylinders, representing a shared environment. Below the cloud, six server racks are shown, representing the community of organizations that share this cloud. Arrows point from the community of organizations up to the community cloud, indicating the flow of data and services.

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.56
------------------	--	-------

56

# PRIVATE CLOUD

- Compute clusters configured as IaaS cloud
- Open source frameworks:
  - Openstack:  
<https://www.openstack.org/>
  - Eucalyptus:  
<https://www.eucalyptus.cloud/>
  - Apache Cloudstack:  
<https://cloudstack.apache.org/>
  - Nimbus:  
<http://www.nimbusproject.org/>
- Various virtualization hypervisors:  
Open source: XEN, KVM    Commercial: VMWare, etc.

organization

private cloud

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.57
------------------	--	-------

57

# HYBRID CLOUD

- Extend private cloud typically with public or community cloud resources
- Cloud bursting:  
Scale beyond one cloud when resource requirements exceed local limitations
- Some resources can remain local for security reasons

organization

public cloud

private cloud

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.58
------------------	--	-------

58

OTHER CLOUDS

- Federated cloud
  - Simply means to aggregate two or more clouds together
  - Hybrid is typically private-public
  - Federated can be public-public, private-private, etc.
  - Also called inter-cloud
- Virtual private cloud
  - Google and Microsoft simply call these virtual networks
  - Ability to interconnect multiple independent subnets of cloud resources together
  - Resources allocated private IPs from individual network subnets can communicate with each other (10.0.1.0/24) and (10.0.2.0/24)
  - Subnets can span multiple availability zones within an AWS region



October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.59

59

TCSS 562  
TERM PROJECT



October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.60

60

TCSS 562 TERM PROJECT

- Build a serverless cloud native application
- Application provides a case study to design trade-offs:
- Projects will compare and contrast one or more trade-offs:
- Service composition
  - Switchboard architecture
    - Address COLD Starts
    - Infrastructure Freeze/Thaw cycle of AWS Lambda (FaaS)
  - Full service isolation, full service aggregation
- Application flow control
- Programming Languages
- Alternate FaaS Platforms
- Data provisioning

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.61

61

EXTRACT TRANSFORM LOAD  
DATA PIPELINE

- Service 1: TRANSFORM
  - Read CSV file, perform some transformations
  - Write out new CSV file
- Service 2: LOAD
  - Read CSV file, load data into relational database
  - Cloud DB (AWS Aurora), or local DB (Derby/SQLite)
    - Derby DB and/or SQLite code examples to be provided in Java

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.62

62

EXTRACT TRANSFORM LOAD  
DATA PIPELINE 2

- Service 3: **EXTRACT**
- Using relational database, apply filter(s) and/or functions to aggregate data to produce sums, totals, averages
- Output aggregations as JSON

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.63

63

SERVICE COMPOSITION

Remote Client

API Gateway

Fine grained services

A	B	C	3 services Full Service Isolation
A	B	C	2 services
A	B	C	2 services
A	B	C	1 service Full Service Aggregation

Other possible compositions: group by library, functional cohesion, etc.

October 16, 2019


TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.64

64



## SWITCH-BOARD ARCHITECTURE



**Single deployment package with consolidated codebase (Java: one JAR file)**

**Entry method contains “switchboard” logic**

**Case statement that route calls to proper service**

**Routing is based on data payload**

**Check if specific parameters exist, route call accordingly**

**Goal: reduce # of COLD starts to improve performance**

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.65

65

## APPLICATION FLOW CONTROL

- **Serverless Computing:**
  - AWS Lambda (FAAS: Function-as-a-Service)
  - Provides HTTP/REST like web services
  - Client/Server paradigm
- **Synchronous web service:**
  - Client calls service
  - Client blocks (freezes) and waits for server to complete call
  - Connection is maintained in the “OPEN” state
  - Problematic if service runtime is long!
    - Connections are notoriously dropped
    - System timeouts reached
- **Client can't do anything while waiting unless using threads**

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.66

66

APPLICATION FLOW CONTROL - 2

- Asynchronous web service
- Client calls service
- Server responds to client with OK message
- Client closes connection
- Server performs the work associated with the service
- Server posts service result in an external data store
  - AWS: S3, SQS (queueing service), SNS (notification service)

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.67

67

APPLICATION FLOW CONTROL - 3

Client flow control

(a)

Microservice as controller

(c)

AWS Step Function

(b)

Asynchronous

(d)

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.68

68

PROGRAMMING LANGUAGE

- Function-as-a-Service platforms support hosting services code in multiple languages
- AWS Lambda- common: Java, Node.js, Python
  - Plus others: Go, PowerShell, C#, and Ruby
- Also Runtime API (“BASH”) which allows deployment of any binary executable in any programming languages
- Jackson D, Clynch G. An Investigation of the Impact of Language Runtime on the Performance and Cost of Serverless Functions. In Proc. Of the 2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion) 2018 Dec 17 (pp. 154-160).
- <http://faculty.washington.edu/wlloyd/courses/tcss562/papers/AnInvestigationOfTheImpactOfLanguageRuntimeOnThePerformanceAndCostOfServerlessFunctions.pdf>

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.69

69

FAAS PLATFORMS

- Many commercial and open source FaaS platforms exist
- TCSS562 projects can choose to compare performance and cost implications of alternate platforms.
- Supported by SAAF:
  - AWS Lambda
  - Google Cloud Functions
  - Azure Functions
  - IBM Cloud Functions

October 16, 2019

TCSS562: Software Engineering for Cloud Computing [Fall 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L7.70

70


# DATA PROVISIONING

- Consider performance and cost implications of the data-tier design for the serverless application
- Use different tools as the relational datastore to support service #2 (LOAD) and service #3 (EXTRACT)
- SQL / Relational:
  - Amazon Aurora (serverless cloud DB), Amazon RDS (cloud DB), DB on a VM (MySQL), DB inside Lambda function (SQLite, Derby)
- NO SQL / Key/Value Store:
  - Dynamo DB, MongoDB, S3

October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.71
------------------	--	-------

71

# QUESTIONS



October 16, 2019	TCSS562: Software Engineering for Cloud Computing [Fall 2019] School of Engineering and Technology, University of Washington - Tacoma	L7.72
------------------	--	-------

72