

TCSS 562:
SOFTWARE ENGINEERING
FOR CLOUD COMPUTING

Cloud Computing
Concepts and Models - III

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



1

OFFICE HOURS – FALL 2025

■ **Thursdays:**

■ 6:00 to 7:00 pm - CP 229 & Zoom

■ **Friday – *** THIS WEEK *****

■ 11:00 am to 12:00 pm – ONLINE via Zoom

■ Or email for appointment

➢ Office Hours set based on Student Demographics survey feedback

➢ * - Friday office hours may be adjusted or canceled due meeting conflicts or other obligations. Adjustments will be announced via Canvas.

October 28, 2025

TCSS462/562: (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.2

2

OBJECTIVES – 10/28

■ **Questions from 10/23**

■ Tutorials Questions

■ Tutorial 5 - Files in S3 and CloudWatch Events

■ Activity 2 Review

■ From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:

- Cloud computing delivery models
- Cloud deployment models

■ AWS Overview and demo

October 28, 2025

TCSS462/562: (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.3

3

ONLINE DAILY FEEDBACK SURVEY

■ Daily Feedback Quiz in Canvas – Take After Each Class

■ Extra Credit for completing

Assignments

Class Activity 1 – Implicit vs. Explicit Parallelism

Available until Oct 13 at 11:59pm | Due Oct 7 at 7:59pm | -15 pts

Tutorial 1 - Linux

Available until Oct 19 at 11:59pm | Due Oct 13 at 11:59pm | -20 pts

TCSS 562 - Online Daily Feedback Survey - 10/5

Available until Oct 18 at 11:59pm | Due Oct 6 at 8:59pm | -15 pts

TCSS 562 - Online Daily Feedback Survey - 9/30

Available until Oct 18 at 11:59pm | Due Oct 4 at 8:59pm | -15 pts

October 28, 2025

TCSS462/562: (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.4

4

TCSS 562 - Online Daily Feedback Survey - 10/5

Started: Oct 7 at 1:13am

Quiz Instructions

Question 1

0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

1 2 3 4 5 6 7 8 9 10

Mostly Review To Me Equal New and Review Mostly New To Me

Question 2

0.5 pts

Please rate the pace of today's class:

1 2 3 4 5 6 7 8 9 10

Slow Just Right Fast

October 28, 2025

TCSS462/562: (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.5

5

MATERIAL / PACE

■ Please classify your perspective on material covered in today's class (43 respondents, 25 in-person, 18 online):

■ 1-mostly review, 5-equal new/review, 10-mostly new

■ **Average – 7.19 (↓ - previous 7.32)**

■ Please rate the pace of today's class:

■ 1-slow, 5-just right, 10-fast

■ **Average – 5.49 (↑ - previous 5.24)**

October 28, 2025

TCSS462/562: (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.6

6

Slides by Wes J. Lloyd

L9.1

QUESTIONS FROM 10/23

▪ **Why does the average \$/hour price within a date range change even though the same region is used (Ohio)**

▪ The average price shows is shown for each availability zone within the region over the time span of graph (3 hrs, 1 day, 1 week, 1 month, 3 months)

Spot Instance pricing history

Your instance type requirements, budget requirements, and application design will determine how to apply the following best practices for your application. To learn more see [Spot Instance Best Practices](#)

Region:

Availability Zone:

Instance type:

Duration:

Price cap:

Filter:

Sort by:

3 hrs 1 day 1 week 1 month 3 months

Time (UTC-07:00)

Price (\$/hr)

Message per hour within date range

Start date:

End date:

Start time:

End time:

Filter:

Sort by:

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.7

7

FEEDBACK - 2

▪ **Least clear point: With more VMs running, the multi-tenancy drops despite best efforts at isolation**

▪ The graph shows increasing multi-tenancy from left to right

▪ more VMs running...

▪ What drops is the performance of our 4 benchmarking applications: sysbench (prime number gen), y-cruncher, pgbench, and iperf

Percentage Change

100%
80%
60%
40%
20%

0 10 20 30 40 50

VM Tenants

sysbench (CPU)

y-cruncher (CPU)

pgbench (CPU + I/O)

iperf (network I/O)

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.8

8

FEEDBACK - 3

▪ **What is "session state" in a client/server application ?**

▪ AWS Lambda functions are considered "stateless"

▪ **What is "stateless" ?**

▪ Each time a client calls an AWS Lambda function, requests are routed to a random function instance (worker) to process the call

▪ If a client makes multiple calls, there is no guarantee it will run in the same function instance each time

▪ AWS Lambda functions feature static global memory, but if client calls do not return to the same function instance, this memory can't be used to store session state

▪ **Key Design Issue With Serverless Applications:**

▪ **How do you persist session state on AWS Lambda ?**

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.9

9

PRACTICE QUESTION 1

▪ Which of the following can lead to performance problems for application hosting on cloud platforms ?

A. Resource sharing/contention

B. Cloud consumer under-provisioning

C. Heterogeneous hardware

D. Cloud provider over-provisioning

E. All of the above

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.10

10

PRACTICE QUESTION 2

▪ Which cloud computing delivery model often requires manual configuration to provide resource elasticity?

A. Platform-as-a-Service

B. Infrastructure-as-a-Service

C. Serverless Database

D. Function-as-a-Service

E. All of the above

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.11

11

PRACTICE QUESTION 3

▪ On Amazon EC2, when using persistent spot requests, what occurs when you intentionally terminate the virtual machine?

A. In addition to the virtual machine being deleted, the persistent spot request is also deleted

B. VM termination is not supported using persistent spot requests

C. Using the AWS management console, the user is prompted to enter a password prior to deletion of the virtual machine

D. After a short delay, a replacement virtual machine is launched to satisfy the persistent spot request

E. The virtual machine is stopped, not terminated, and can be later resumed without loss of data on the disk

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.12

12

Slides by Wes J. Lloyd

L9.2

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions**
- Tutorial 5 - Files in S3 and CloudWatch Events
- Activity 2 Review
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models
- AWS Overview and demo

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.13

13

TUTORIAL 3 – OCT 30 (TEAMS OF 2)

- Best Practices for Working with Virtual Machines on Amazon EC2
- https://faculty.washington.edu/wlloyd/courses/tcss562/tutorials/TCSS462_562_f2025_tutorial_3.pdf
- Creating a spot VM
- Creating an image from a running VM
- Persistent spot request
- Stopping (pausing) VMs
- EBS volume types
- Ephemeral disks (local disks)
- Mounting and formatting a disk
- Disk performance testing with Bonnie++
- Cost Saving Best Practices

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.14

14

TUTORIAL 4 – NOV 7

- Introduction to AWS Lambda with the Serverless Application Analytics Framework (SAAF)
https://faculty.washington.edu/wlloyd/courses/tcss562/tutorials/TCSS462_562_f2025_tutorial_4.pdf
- Set up Java development environment
- Introduction to Maven build files for Java
- Create and Deploy "hello" Java AWS Lambda Function
- Create API Gateway REST endpoint
- Sequential testing of "hello" AWS Lambda Function
 - API Gateway endpoint, AWS Lambda CLI Function Invocation, AWS Function URL
- Profiling function performance with SAAF
- Concurrent function testing with faas_runner
- Performance analysis using faas_runner reports
- Two function pipeline development task: Caesar Cipher

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.15

15

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions
- Tutorial 5 - Files In S3 and CloudWatch Events**
- Activity 2 Review
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models
- AWS Overview and demo

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.16

16

TUTORIAL 5 – TO BE POSTED

- Introduction to Lambda II: Working with Files in S3 and CloudWatch Events
- Customize the Request object (add getters/setters)
 - Why do this instead of HashMap ?
- Import dependencies (jar files) into project for AWS S3
- Create an S3 Bucket
- Give your Lambda function(s) permission to work with S3
- Write to the CloudWatch logs
- Use of CloudTrail to generate S3 events
- Creating CloudWatch rule to capture events from CloudTrail
- Have the CloudWatch rule trigger a target Lambda function with a static JSON input object (hard-coded filename)
- Optional:** for the S3 PutObject event, dynamically extract the name of the file put to the S3 bucket for processing

October 27, 2022

TCSS462/562: Software Engineering for Cloud Computing [Fall 2022]
School of Engineering and Technology, University of Washington - Tacoma

L9.17

17

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions
- Tutorial 5 - Files in S3 and CloudWatch Events
- Activity 2 Review**
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models
- AWS Overview and demo

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.18

18

CLOUD COMPUTING:
CONCEPTS AND MODELS



October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.19

19

CATCH UP FROM 10/23

- Questions from 10/16
- Tutorials Questions
- Tutorial 4 – Intro to FaaS – AWS Lambda
- Background on AWS Lambda for the Term Project - II
- From: **Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:**
 - Roles and boundaries
 - **Cloud characteristics**
 - Cloud delivery models
 - Cloud deployment models
- Team Planning - Breakout Rooms

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.20

20

CLOUD CHARACTERISTICS

- On-demand usage
- Ubiquitous access
- Multitenancy (resource pooling)
- Elasticity
- **Measured usage**
- Resiliency

■ Assessing these features helps measure the value offered by a given cloud service or platform

October 17, 2024

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L7.21

21

MEASURED USAGE

- Cloud platform tracks usage of IT resources
- For billing purposes
- Enables charging only for IT resources actually used
- Can be time-based (millisec, second, minute, hour, day)
 - Granularity is increasing...
- Can be throughput-based (data transfer: MB/sec, GB/sec)
- Can be resource/reservation based (vCPU/hr, GB/hr)

■ Not all measurements are for billing

■ Some measurements can support auto-scaling

■ For example CPU utilization


October 17, 2024

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L7.22

22

EC2 CLOUDWATCH METRICS



October 17, 2024

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L7.23

23

EC2 CLOUDWATCH METRICS



October 17, 2024

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L7.24

24

RESILIENCY

- Distributed redundancy across physical locations (regions on AWS)
- Used to improve reliability and availability of cloud-hosted applications
- Very much an engineering problem
- No "resiliency-as-a-service" for user deployed apps
- Unique characteristics of user applications make a one-size fits all service solution challenging

October 17, 2024

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L7.25

25

Elasticity is often provided using threshold based scaling. When can threshold based scaling (i.e. CPU utilization > 80%) under or over provision resources?

When the application is primarily I/O bound, a CPU threshold may never be met, or be met too late to scale up.

When the current resource utilization does not reflect future system demand.

When the current resource utilization (e.g. CPU) is temporarily increased as a result of external factors (e.g. resource contention from other tasks) that does not correlate to system demand.

When an application will soon complete a parallel phase, before executing a largely sequential phase

All of the above

October 24, 2018

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L10.28

26

When poll is active, respond at polllev.com/wesleylloyd641

Text **WESLEYLLOYD641** to **22333** once to join

The scaling threshold of "when CPU utilization > 80% scale up", is:

An application specific threshold

An application agnostic threshold

Start the presentation to see live content. For screen share software, share the entire screen. Get help at polllev.com/app

27

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions
- Tutorial 5 - Files in S3 and CloudWatch Events
- Activity 2 Review
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models
- AWS Overview and demo

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.28

28

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.29

29

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS) delivery model
- Virtualization is a key-enabling technology of IaaS cloud
- Uses virtual machines to deliver cloud resources to end users

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.30

30

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS) delivery model

Virtualization is key to sharing powerful servers among users by running *many* isolated private virtual computers known as virtual machines (VMs)
...VMs are the basis of cloud v1.0

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.31

31

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS) delivery model

Virtual Machines are the building blocks for "Cloud Service Delivery Models"
They are the "vehicles" used to deliver compute resources to end users...
cloud 1.0



October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.32

32

CLOUD COMPUTING 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?
- As users, how do we minimize hosting costs?
 - How do we estimate hosting costs?




October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.33

33

CLASSIC CLOUD COMPUTING DELIVERY MODELS




October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.34

34

CLASSIC CLOUD COMPUTING DELIVERY MODELS



October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.35

35

EXAMPLE CLOUD SERVICES

 SAAS Software as a Service	 PAAS Platform as a Service	 IAAS Infrastructure as a Service
Email CRM Collaborative ERP	Application Development Decision Support Web Streaming	Caching Legacy Networking Security File Technical System Mgmt
CONSUME	BUILD ON IT	MIGRATE TO IT

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.36

36

END USER APPLICATIONS

Many different “cloud” providers (especially SaaS)

Many cloud providers are also cloud consumers

Software-as-a-Service

Infrastructure-as-a-Service

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.37

37

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**: Scripts to specify creation of cloud infrastructures using JSON/YAML for app deployment

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.38

38

IAAS: VIRTUAL MACHINE PLACEMENT IN THE CLOUD

Number of VMs
Objective Constraints

Requests

User

Use

Placement Algorithm

VM type 1

VM type 2

...

VM type n

Virtual Machine Manager

Cloud Broker

Cloud 1

Cloud 1

...

Cloud n

Controlled by the cloud provider

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.39

39

COMPONENT COMPOSITION

- Cloud provider maps VMs to physical servers
- User controls mapping of services to VM images
 - Should components be separated (isolated) ?
 - Should components be combined ?
 - M – modeling web service
 - D – relational database
 - F – file server
 - L – logging server

SC1

SC2

SC3

SC4

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.40

40

SC1

SC2

SC3

SC4

SC5

SC6

SC7

SC8

SC9

SC10

SC11

SC12

SC13

SC14

SC15

M: Tomcat Application Server
D: PostgreSQL DB
F: nginx file server
L: Log server (Codebeamer)

41

41

SC1

SC2

SC3

SC4

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

SC15


M: Tomcat Application Server
D: PostgreSQL DB
F: nginx file server
L: Log server (Codebeamer)

42


42




43




44




45



46



47



48

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.49


49

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)



The diagram illustrates the Platform-as-a-Service (PaaS) architecture. On the left, several client devices (laptops) are shown. Arrows from these clients point to a central 'Load Balancer' box. From the Load Balancer, arrows point to a group of 'Servers' on the right. The servers are represented by server rack icons.

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.50

50

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 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.51

51

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.52

52

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 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.53

53



The diagram illustrates the multi-layered architecture of cloud computing. It is organized into three horizontal layers. The top layer, labeled 'SaaS' on the left, contains three yellow circles representing 'Cloud Service A', 'Cloud Service B', and 'Cloud Service C'. The middle layer, labeled 'PaaS' on the left, contains three boxes representing 'Ready-Made Environment A', 'Ready-Made Environment B', and 'Ready-Made Environment C'. The bottom layer, labeled 'IaaS' on the left, contains two boxes representing 'Virtual Server A' and 'Virtual Server B', which are shown connected to a single 'Physical Server A' at the base. Dashed lines indicate the flow of services and dependencies between these layers.

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.54

54

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 23, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.55

55

SERVERLESS COMPUTING

Introducing Cloud 2.0

Serverless Computing

Deploy Applications Without Fiddling With Servers

October 23, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.56

56

SERVERLESS COMPUTING

How should I tell if a server has been compromised?

How can I tell if a server has been compromised?

How can I increase utilization of my servers?

Which OS should my servers run?

How much remaining capacity do my servers have?

How should I implement dynamic configuration changes on my servers?

How will I keep my server OS patched?

How can I control access from my servers?

How will new code be deployed to my servers?

What size servers are right for my budget?

When should I decide to scale up my servers?

Which packages should be baked into my server images?

How will the application handle server hardware failure?

Which users should have access to my servers?

Should I tune OS settings to optimize my application?

How many users create too much load for my servers?

How many servers should I budget for?

When should I decide to scale out my servers?

What size server is right for my performance?

Servers

(AHHHHHHHHH!)

October 23, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.57

57

SERVERLESS COMPUTING

What is serverless?

Build and run applications without thinking about servers

October 23, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.58

58

SERVERLESS COMPUTING - 2

Evolving to serverless

October 23, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.59

59

SERVERLESS COMPUTING

Pay only for CPU/memory utilization

High Availability

Fault Tolerance

Infrastructure Elasticity

No Setup

Function-as-a-Service (FAAS)

60

SERVERLESS COMPUTING

Why Serverless Computing?

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

...they are built into the platform

61

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.62

62

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 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.63

63

FAAS PLATFORMS

The diagram shows a cloud shape containing two groups of boxes. The 'Commercial' group includes AWS Lambda, Azure Functions, IBM Cloud Functions, and Google Cloud Functions. The 'Open Source' group includes Apache OpenWhisk and Fn (Oracle).

64

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 ranging 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

October 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.65

65

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 23, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.66

66

FAAS PLATFORMS - 3

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

October 23, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.67

67

CLOUD NATIVE SOFTWARE ARCHITECTURE

- Every service with a different pricing model

Example: Weather Application

Front-end code for weather app hosted in S3

User clicks on link to get local weather information

API GATEWAY

Lambda is triggered

35° C

Lambda runs code to retrieve local weather information and returns data back to user

DYNAMODB

October 23, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.68

68

IAAS BILLING MODELS

- Virtual machines as-a-service at ¢ per hour
- No premium to scale:
$$\begin{matrix} 1000 \text{ computers} & @ & 1 \text{ hour} \\ = & & 1 \text{ computer} & @ & 1000 \text{ hours} \end{matrix}$$
- 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 →

October 23, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.69

69

PRICING OBFUSCATION

- VM pricing:** hourly rental pricing, billed to nearest second is intuitive...
- FaaS pricing:** non-intuitive pricing policies
- 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
\$0.00000208 to rent 128MB / 100-ms
\$0.0001667 GB / second

October 23, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.70

70

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)
- Serverless Computing:**
 - Function-as-a-Service (FaaS)**
 - Container-as-a-Service (CaaS)
- Other Delivery Models

October 22, 2024

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.71

71

WEBSERVICE HOSTING EXAMPLE

- ON AWS Lambda**
- Each service call: 100% of 2 CPU-cores
100% of 4GB of memory
- Workload: uses 2 continuous threads
- Duration: 1 month (30.41667 days)
- ON AWS EC2:** Amazon EC2 c5.large 2-vCPU VM x 4GB
- c5.large: 8.5¢/hour, 24 hrs/day x 30.41667 days
- Hosting cost: \$62.05/month
- How much would hosting this workload cost on AWS Lambda?**

October 22, 2024

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.72

72

PRICING OBFUSCATION

Assume 1 month = 30.41667 days (365d / 12)

Worst-case FaaS scenario = ~2.72x !

AWS EC2: \$62.05

AWS Lambda: \$168.91

Break Even: 3,702,459 GB-sec

@4GB ~10.71 days

BREAK-EVEN POINT: \$62.05 - \$0.33 (calls) = \$61.72

\$61.72 / .00001667 GB-sec = ~3,702,459 GB-sec-mon/4GB/call=

~925,614 sec or ~10.71 days

Point at which using FaaS costs the same as IaaS

73

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?

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.74

74

FAAS CHALLENGES

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?

Performance – what will it be?

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.75

75

VENDOR ARCHITECTURAL LOCK-IN

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

Example: Weather Application

S3

Client

API GATEWAY

35° C

Lambda is triggered

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 runs code to retrieve local weather information and returns data back to user

Images credit: aws.amazon.com

Increased dependencies → increased hosting costs

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.76

76

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 400,000 GB-sec/month → FREE

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

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.77

77

MEMORY RESERVATION QUESTION...

Lambda memory reserved for functions

UI provides text box formerly "slider bar" to set function's memory

Resource capacity (CPU, disk, network) coupled to slider bar:
"every doubling of memory, doubles CPU..."

But how much memory do FaaS functions 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

Memory (MB) info

Your function is allocated CPU proportional to the memory configured.

1024 MB

Set memory to between 128 MB and 10240 MB

October 28, 2025


TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.78


78

Slides by Wes J. Lloyd


L9.13




79




80




81



82



83



84

AWS LAMBDA PERFORMANCE VARIATION - 2

- NLP use case: Less performance variance using ARM-based CPUs (less resource contention), and w/o hyperthreading

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma L9.85

85

FUNCTION-AS-A-SERVICE

AWS
Lambda
Demo

86

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma L9.87

87

CONTAINER-AS-A-SERVICE

- Cloud service model for deploying application containers (e.g. Docker containers) to the cloud
- Deploy containers without worrying about managing infrastructure:
 - Servers (virtual machines)
 - 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
 - Google Cloud Run
 - Azure Container Instances

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma L9.88

88

CLOUD COMPUTING DELIVERY MODELS

- Infrastructure-as-a-Service (IaaS)
- Platform-as-a-Service (PaaS)
- Software-as-a-Service (SaaS)

Serverless Computing:

- Function-as-a-Service (FaaS)
- Container-as-a-Service (CaaS)
- Other Delivery Models

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma L9.89

89

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 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma L9.90

90

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions
- Tutorial 5 - Files in S3 and CloudWatch Events
- Activity 2 Review
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models**
- AWS Overview and demo

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.91

91

CLOUD DEPLOYMENT MODELS

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


October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.92

92

PUBLIC CLOUDS



The diagram illustrates public clouds as a collection of services from different providers. At the top, there are seven cloud icons, each labeled with a company name: Google, Salesforce, Microsoft, Yahoo, Amazon, Zoho, and Rackspace. Below these, three arrows point upwards towards a central icon of a city skyline labeled 'organizations', indicating that these providers serve various organizations.

October 28, 2025


TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.93

93

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 shows a community cloud as a shared resource. A large cloud icon contains several server racks and storage symbols. Below the cloud, three arrows point upwards towards a central icon of a city skyline labeled 'community of organizations', indicating that the cloud is shared among a specific group of organizations.

October 28, 2025


TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.94

94

PRIVATE CLOUD

- Compute clusters configured as IaaS cloud
- Open source software
 - Eucalyptus
 - Openstack
 - Apache Cloudstack
 - Nimbus
- Virtualization: XEN, KVM, ...



The diagram depicts a private cloud setup. An 'organization' icon (city skyline) is connected to a 'cloud service consumer' icon (server rack). This consumer is linked to a 'cloud service' icon (server rack) which is located inside a cloud icon labeled 'private cloud'.

October 28, 2025


TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.95

95

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



The diagram illustrates a hybrid cloud model. It shows a 'public cloud' (cloud icon) and a 'private cloud' (cloud icon) connected by a double-headed arrow. Both clouds are connected to a central 'organization' icon (city skyline). The public cloud contains a 'cloud service' icon, and the private cloud contains a 'cloud service' icon.

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.96

96

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 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.97

97

WE WILL RETURN AT
4:50 PM



98

OBJECTIVES – 10/28

- Questions from 10/23
- Tutorials Questions
- Tutorial 5 - Files in S3 and CloudWatch Events
- Activity 2 Review
- From: Cloud Computing Concepts, Technology & Architecture: Chapter 4: Cloud Computing Concepts and Models:
 - Cloud computing delivery models
 - Cloud deployment models
- AWS Overview and demo**

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.99

99

AWS OVERVIEW
AND DEMO



100

ONLINE CLOUD TUTORIALS

- From the eScience Institute @ UW Seattle:
<https://escience.washington.edu/>
- Online cloud workshops
- Introduction to AWS, Azure, and Google Cloud
- Task: Deploying a Python DJANGO web application
- Self-guided workshop materials available online:
<https://cloudmaven.github.io/documentation/>
- AWS Educate provides access to many online tutorials / learning resources:
<https://aws.amazon.com/education/awseducate/>

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.101

101

LIST OF TOPICS

- AWS Management Console
- Elastic Compute Cloud (EC2)
- Instance Storage: Virtual Disks on VMs
- Elastic Block Store: Virtual Disks on VMs
- Elastic File System (EFS)
- Amazon Machine Images (AMIs)
- EC2 Paravirtualization
- EC2 Full Virtualization (hvm)
- EC2 Virtualization Evolution

- (VM) Instance Actions
- EC2 Networking
- EC2 Instance Metadata Service
- Simple Storage Service (S3)
- AWS Command Line Interface (CLI)
- Legacy / Service Specific CLIs
- AMI Tools
- Signing Certificates
- Backing up live disks
- Cost Savings Measures


October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.10
2

102

AWS MANAGEMENT CONSOLE



October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.103

103

AWS EC2

- **Elastic Compute Cloud**
- Instance types: <https://ec2instances.info>
 - **On demand Instance** – full price
 - **Reserved Instance** – contract based where customer guarantees VM rental for a fixed period of time (e.g. 1 year, 3 years, etc.)
 Deeper discounts with longer term commitments
 - **Spot Instance** – portion of cloud capacity reserved for low cost instances, when demand exceeds supply instances are randomly terminated with 2 minute warning
 - Users can make diverse VM requests using different types, zones, regions, etc. to minimize instance terminations
 - Developers can design for failure because often only 1 or 2 VMs in a cluster fail at any given time. They then need to be replaced.
 - **Dedicated host** – reserved private HW (server)
 - Instance families –
 - General, compute-optimized, memory-optimized, GPU, etc.

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.104

104

AWS EC2 - 2

- **Storage types**
 - **Instance storage - ephemeral storage**
 - Temporary disk volumes stored on disks local to the VM
 - Evolution: physical hard disk drives (HDDs)
 - Solid state drives (SSDs)
 - Non-volatile memory express (NVMe) drives (closer to DRAM speed)
 - **EBS - Elastic block store**
 - Remotely hosted disk volumes
 - **EFS - Elastic file system**
 - Shared file system based on network file system
 - VMs, Lambdas, Containers mount/interact with shared file system
 - Somewhat expensive

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.105

105

INSTANCE STORAGE

- Also called ephemeral storage
- Persisted using images saved to S3 (simple storage service)
 - ~2.3¢ per GB/month on S3
 - 5GB of free tier storage space on S3
- Requires “burning” an image
- Multi-step process:
 - Create image files
 - Upload chunks to S3
 - Register image
- Launching a VM
 - Requires downloading image components from S3, reassembling them... is potentially slow
- VMs with instance store backed root volumes not pause-able
- Historically root volume limited to 10-GB max – **faster Imaging...**

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.106

106

ELASTIC BLOCK STORE

- EBS provides 1 drive to 1 virtual machine (**1 : 1 (not shared)**)
- EBS cost model is different than instance storage (uses S3)
 - ~10¢ per GB/month for General Purpose Storage (GP2)
 - ~8¢ per GB/month for General Purpose Storage (GP3)
 - 30GB of free tier storage space
- EBS provides “live” mountable volumes
 - Listed under volumes
 - **Data volumes:** can be mounted/unmounted to any VM, dynamically at any time
 - **Root volumes:** hosts OS files and acts as a boot device for VM
 - In Linux drives are linked to a mount point “directory”
- Snapshots back up EBS volume data to S3
 - Enables replication (required for horizontal scaling)
 - EBS volumes not actively used should be snapshotted, and deleted to save EBS costs...

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.107

107

EBS VOLUME TYPES - 2

- **Metric: I/O Operations per Second (IOPS)**
- **General Purpose 2 (GP2)**
 - 3 IOPS per GB, min 100 IOPS (<34GB), max of 16,000 IOPS
 - 250MB/sec throughput per volume
- **General Purpose 3 (GP3 – new Dec 2020)**
 - Max 16,000 IOPS, Default 3,000 IOPS
 - GP2 requires creating a 1TB volume to obtain 3,000 IOPS
 - GP3 all volumes start at 3000 IOPS and 125 MB/s throughput
 - 1000 additional IOPS beyond 3000 is \$5/month up to 16000 IOPS
 - 125 MB/s additional throughput is \$5/month up to 1000 MB/s throughput

October 28, 2025 TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma L9.108

108

EBS VOLUME TYPES - 3

- **Provisioned IOPS (IO1)**
 - Legacy, associated with GP2
 - Allows user to create custom disk volumes where they pay for a specified IOPS and throughput
 - 32,000 IOPS, and 500 MB/sec throughput per volume MAX
- **Throughput Optimized HDD (ST1)**
 - Up to 500 MB/sec throughput
 - 4.5 ¢ per GB/month
- **Cold HDD (SC1)**
 - Up to 250 MB/sec throughput
 - 2.5 ¢ per GB/month
- **Magnetic**
 - Up to 90 MB/sec throughput per volume
 - 5 ¢ per GB/month

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.109

109

ELASTIC FILE SYSTEM (EFS)

- EFS provides 1 volume to many client (1 : n) **shared storage**
- Network file system (based on NFSv4 protocol)
- Shared file system for EC2, Fargate/ECS, Lambda
- Enables mounting (sharing) the same disk "volume" for R/W access across multiple instances at the same time
- Different performance and limitations vs. EBS/Instance store
- Implementation uses abstracted EC2 instances
- ~ 30 ¢ per GB/month storage – **default burstable throughput**
- **Throughput modes:**
 - Can modify modes only once every 24 hours
- **Burstable Throughput Model:**
 - Baseline – 50kb/sec per GB
 - Burst – 100MB/sec per GB (for volumes sized 10GB to 1024 GB)
 - Credits - .72 minutes/day per GB

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.110

110

ELASTIC FILE SYSTEM (EFS) - 2

- **Burstable Throughput Rates** *Information subject to revision*
 - Throughput rates: baseline vs burst
 - Credit model for bursting: maximum burst per day

File System Size (GiB)	Baseline Aggregate Throughput (MiB/s)	Burst Aggregate Throughput (MiB/s)	Maximum Burst Duration (Min/Day)	% of Time File System Can Burst (Per Day)
10	0.5	100	7.2	0.5%
256	12.5	100	180	12.5%
512	25.0	100	360	25.0%
1024	50.0	100	720	50.0%
1536	75.0	150	720	50.0%
2048	100.0	200	720	50.0%
3072	150.0	300	720	50.0%
4096	200.0	400	720	50.0%

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.111

111

ELASTIC FILE SYSTEM (EFS) - 3

- **Throughput Models** *Information subject to revision*
- **Provisioned Throughput Model**
- For applications with:
 - high performance requirements, but low storage requirements
- Get high levels of performance w/o overprovisioning capacity
- **\$6 MB/s-Month (Virginia Region)**
 - Default is 50kb/sec for 1 GB, .05 MB/s = 30 ¢ per GB/month
- If file system metered size has higher baseline rate based on size, file system follows default Amazon EFS Bursting Throughput model
 - No charges for Provisioned Throughput below file system's entitlement in Bursting Throughput mode
 - Throughput entitlement = 50kb/sec per GB

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.112

112

ELASTIC FILE SYSTEM (EFS) - 4

Performance Comparison, Amazon EFS and Amazon EBS *Information subject to revision*

	Amazon EFS	Amazon EBS Provisioned IOPS
Per-operation latency	Low, consistent latency.	Lowest, consistent latency.
Throughput scale	10+ GB per second.	Up to 2 GB per second.

Storage Characteristics Comparison, Amazon EFS and Amazon EBS

	Amazon EFS	Amazon EBS Provisioned IOPS
Availability and durability	Data is stored redundantly across multiple AZs.	Data is stored redundantly in a single AZ.
Access	Up to thousands of Amazon EC2 instances, from multiple AZs, can connect concurrently to a file system.	A single Amazon EC2 instance in a single AZ can connect to a file system.
Use cases	Big data and analytics, media processing workflows, content management, web serving, and home directories.	Boot volumes, transactional and NoSQL databases, data warehousing, and ETL.

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.113

113

AMAZON MACHINE IMAGES

- AMIs
- Unique for the operating system (root device image)
- Two types
 - Instance store
 - Elastic block store (EBS)
- Deleting requires multiple steps
 - Deregister AMI
 - Delete associated data - (files in S3)
- Forgetting both steps leads to costly "orphaned" data
 - No way to instantiate a VM from deregistered AMIs
 - Data still in S3 resulting in charges

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.114

114

EC2 VIRTUALIZATION - PARAVIRTUAL

- 1st, 2nd, 3rd, 4th generation → XEN-based
- 5th generation Instances → AWS Nitro virtualization
- XEN - two virtualization modes
- XEN Paravirtualization “paravirtual”
 - 10GB Amazon Machine Image – base image size limit
 - Addressed poor performance of old XEN HVM mode
 - I/O performed using special XEN kernel with XEN paravirtual mode optimizations for better performance
 - Requires OS to have an available paravirtual kernel
 - PV VMs: will use common **AKI** files on AWS – **Amazon kernel Image(s)**
 - Look for common identifiers

October 28, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.115

115

EC2 VIRTUALIZATION - HVM

- XEN HVM mode
 - Full virtualization – no special OS kernel required
 - Computer entirely simulated
 - MS Windows runs in “hvm” mode
 - Allows work around: 10GB instance store root volume limit
 - Kernel is on the root volume (under /boot)
 - No AKIs (kernel images)
 - Commonly used today (*EBS-backed instances*)

October 28, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.116

116

EC2 VIRTUALIZATION - NITRO

- Nitro based on Kernel-based-virtual-machines
 - Stripped down version of Linux KVM hypervisor
 - Uses KVM core kernel module
 - I/O access has a direct path to the device
- Goal: provide indistinguishable performance from bare metal

October 28, 2025


TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.117

117

EVOLUTION OF AWS VIRTUALIZATION

From: <http://www.brendangregg.com/blog/2017-11-29/aws-ec2-virtualization-2017.html>



October 28, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.118

118

INSTANCE ACTIONS

- Stop
 - Costs of “pausing” an instance
- Terminate
- Reboot
- Image management
- Creating an image
 - EBS (snapshot)
- Bundle image
 - Instance-store

October 28, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.119

119

EC2 INSTANCE: NETWORK ACCESS

- Public IP address
- Elastic IPs
 - Costs: in-use FREE, not in-use ~12 \$/day
 - Not in-use (e.g. “paused” EBS-backed instances)
- Security groups
 - E.g. firewall
- Identity access management (IAM)
 - AWS accounts, groups
- VPC / Subnet / Internet Gateway / Router
- NAT-Gateway

October 28, 2025

TCSS462/562:Software Engineering for Cloud Computing [Fall 2025]
 School of Engineering and Technology, University of Washington - Tacoma

L9.120

120

SIMPLE VPC

- Recommended when using Amazon EC2

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.121

121

VPC SPANNING AVAILABILITY ZONES

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.124

122

INSPECTING INSTANCE INFORMATION

- EC2 VMs run a local metadata service
- Can query instance metadata to self discover cloud configuration attributes
- Find your instance ID:

```
curl http://169.254.169.254/
```

```
curl http://169.254.169.254/latest/
```

```
curl http://169.254.169.254/latest/meta-data/
```

```
curl http://169.254.169.254/latest/meta-data/instance-id
```

```
; echo
```
- `ec2-get-info` command
- Python API that provides easy/formatted access to metadata

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.123

123

SIMPLE STORAGE SERVICE (S3)

- Key-value blob storage
- What is the difference vs. key-value stores (NoSQL DB)?
- Can mount an S3 bucket as a volume in Linux
 - Supports common file-system operations
- Provides eventual consistency
- Can store Lambda function state for life of container.

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.124

124

AWS CLI

- Launch Ubuntu 16.04 VM
 - Instances | Launch Instance
- Install the general AWS CLI
 - `sudo apt install awscli`
- Create config file

```
[default]  
aws_access_key_id = <access key id>  
aws_secret_access_key = <secret access key>  
region = us-east-1
```

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.125

125

AWS CLI - 2

- Creating access keys: IAM | Users | Security Credentials | Access Keys | Create Access Keys

October 28, 2025

TCSS462/562: Software Engineering for Cloud Computing [Fall 2025]
School of Engineering and Technology, University of Washington - Tacoma

L9.126

126

AWS CLI - 3

- Export the config file
 - Add to /home/ubuntu/.bashrc

```
export AWS_CONFIG_FILE=$HOME/.aws/config
```
- Try some commands:
 - aws help
 - aws command help
 - aws ec2 help
 - aws ec2 describes-instances --output text
 - aws ec2 describe-instances --output json
 - aws s3 ls
 - aws s3 ls vmscaleruw

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.127

127

LEGACY / SERVICE SPECIFIC CLI(S)

- sudo apt install ec2-api-tools
- Provides more concise output
- Additional functionality
- Define variables in .bashrc or another sourced script:
 - export AWS_ACCESS_KEY={your access key}
 - export AWS_SECRET_KEY={your secret key}
- ec2-describe-instances
- ec2-run-instances
- ec2-request-spot-instances
- EC2 management from Java:
 - <http://docs.aws.amazon.com/AWSJavaSDK/latest/javadoc/index.html>
- Some AWS services have separate CLI installable by package

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.128

128

AMI TOOLS

- Amazon Machine Images tools
- For working with disk volumes
- Can create live copies of any disk volume
 - Your local laptop, ec2 root volume (EBS), ec2 ephemeral disk
- Installation:
 - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ami-tools-commands.html>
- AMI tools reference:
 - <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ami-tools-commands.html>
- Some functions may require private key & certificate files

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.129

129

PRIVATE KEY AND CERTIFICATE FILE

- Install openssl package on VM
- # generate private key file

```
$openssl genrsa 2048 > mykey.pk
```
- # generate signing certificate file

```
$openssl req -new -x509 -nodes -sha256 -days 36500 -key mykey.pk -outform PEM -out signing.cert
```
- Add signing.cert to IAM | Users | Security Credentials |
 - new signing certificate --
- From: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/setup-ami-tools.html?icmpid=docs_iam_console#ami-tools-create-certificate

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.130

130

PRIVATE KEY, CERTIFICATE FILE

- These files, combined with your AWS_ACCESS_KEY and AWS_SECRET_KEY and AWS_ACCOUNT_ID enable you to publish new images from the CLI
- Objective:
 - Configure VM with software stack
 - Burn new image for VM replication (**horizontal scaling**)
- An alternative to bundling volumes and storing in S3 is to use a containerization tool such as Docker. . .
- Create image script . . .

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.131

131

SCRIPT: CREATE A NEW INSTANCE STORE IMAGE FROM LIVE DISK VOLUME

```
image=$1
echo "Burn image $image"
echo "$image" > image.id
mkdir /mnt/tmp
AWS_KEY_DIR=/home/ubuntu/.aws
export EC2_URL=http://ec2.amazonaws.com
export S3_URL=https://s3.amazonaws.com
export EC2_PRIVATE_KEY=${AWS_KEY_DIR}/mykey.pk
export EC2_CERT=${AWS_KEY_DIR}/signing.cert
export AWS_USER_ID={your account id}
export AWS_ACCESS_KEY={your aws access key}
export AWS_SECRET_KEY={your aws secret key}
ec2-bundle-vol -s 5000 -u ${AWS_USER_ID} -c ${EC2_CERT} -k ${EC2_PRIVATE_KEY}
--ec2cert /etc/ec2/ami-tools/cert-ec2.pem --no-inherit -r x86_64 -p $image -i /etc/ec2/ami-tools/cert-ec2.pem
cd /tmp
ec2-upload-bundle -b tcss562 -w $image.manifest.xml -a ${AWS_ACCESS_KEY} -s ${AWS_SECRET_KEY} --url http://s3.amazonaws.com --location US
ec2-register tcss562/$image.manifest.xml --region us-east-1 --kernel aki-88aa75e1
```

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.132

132

COST SAVINGS MEASURES

▪ **From Tutorial 3:**

▪ **#1:** ALWAYS USE SPOT INSTANCES FOR COURSE/RESEARCH RELATED PROJECTS

▪ **#2:** NEVER LEAVE AN EBS VOLUME IN YOUR ACCOUNT THAT IS NOT ATTACHED TO A RUNNING VM

▪ **#3:** BE CAREFUL USING PERSISTENT REQUESTS FOR SPOT INSTANCES

▪ **#4:** TO SAVE/PERSIST DATA, USE EBS SNAPSHOTS AND THEN

▪ **#5:** DELETE EBS VOLUMES FOR TERMINATED EC2 INSTANCES.

▪ **#6:** UNUSED SNAPSHOTS AND UNUSED EBS VOLUMES SHOULD BE PROMPTLY DELETED !!

▪ **#7:** USE PERSISTENT SPOT REQUESTS AND THE "STOP" FEATURE TO PAUSE VMS DURING SHORT BREAKS

October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.133

133

QUESTIONS



October 28, 2025

TCSS462/562 (Software Engineering for) Cloud Computing (Fall 2025)
School of Engineering and Technology, University of Washington - Tacoma

L9.134

134