

## TCSS 562: SOFTWARE ENGINEERING FOR CLOUD COMPUTING

### Cloud Computing: AWS Demo

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



## FEEDBACK FROM 10/24

- Which Operating System should I use for TCSS562 projects?
  - Entirely possible to complete tutorials on Windows / Linux / Mac, but the tutorials are written for Ubuntu / Linux.
  - Development operating systems (OS) are at the discretion of the student
  - Employers generally do not dictate a development OS, but are expected to know how to use at least one well
    - See a variety of laptops: MAC, Windows, Linux

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.2

## FEEDBACK - 2

- Tradeoffs of sticking with Windows/Mac:
  - Gains a practicum on interpreting how to perform Linux operations on a Windows or Mac system.
  - Large volume of documentation online is geared towards Linux (e.g. StackOverflow.com)
  - Must reinterpret instructions to Windows/Mac
  - Ability to translate can be beneficial in these scenarios
- Tradeoffs If migrating to Linux from Windows/Mac:
  - Gain valuable skills working with Linux (Ubuntu)
  - Linux is the predominant cloud operating system for virtualization
  - Save time from translating instructions to other platforms potentially making development more productive

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.3

## FEEDBACK - 3

- Recall article on predominance of Linux on the Azure Cloud
- <https://www.zdnet.com/article/linux-now-dominates-azure/>
- The purpose of Tutorials 1 & 2 was to introduce a Linux environment.
- Licensing costs of commercial operating systems such as Windows and Mac make them less attractive for backend development because high licensing costs can severely limit dynamic scaling.

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.4

## FEEDBACK - 4

- Is there any difference in communication (network) latency between a DBaaS (Amazon Relational Database Service – RDS) and?
  - An application server (e.g. Apache Tomcat) hosted on an EC2 instance
  - Application code deployed to AWS Lambda
- Think of AWS Lambda as a c4.large EC2 instance
- Memory setting controls CPU power and network throughput

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.5

## FEEDBACK - 5

- Location? Invocation? of containers when AWS Lambda is called using bash script
- Lambda code had access to 512MB temporary file system
  - Directory is "/tmp"
  - Can read/write local files here
- To identify unique containers, write a file to /tmp with a UUID
- If UUID-file already exists: **detect a reused container**
- If UUID-file missing: **detect a new container**
- VM is identified by the btime from /proc/stat
- Lambda runs AWS Linux, based on RedHat Linux (CentOS)

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.6

FEEDBACK - 6

- What types of questions we will have on midterm?
- I plan to review from now ...

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L10.7

OBJECTIVES


- Tutorial 4
- Tutorial 5
- Midterm 11/7
- Term Project Check-in
- AWS Demo

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L10.8

AWS DEMO



October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

L5.9

CLOUD 101 WORKSHOP

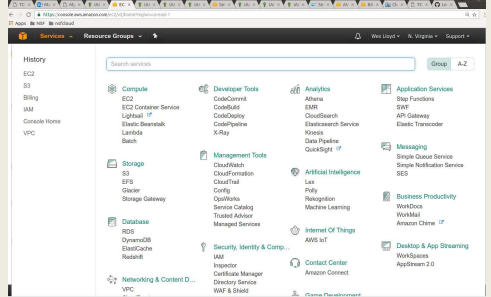
- From the eScience Institute @ UW Seattle:
  - <https://escience.washington.edu/>
  - Offers 1-day cloud workshops
- Introduction to AWS, Azure, and Google Cloud
- Task: Deploying a Python DJANGO web application
- Workshop materials available online:
  - [https://cloudmaven.github.io/documentation/rc\\_cloud101\\_immersion.html](https://cloudmaven.github.io/documentation/rc_cloud101_immersion.html)

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

L5.10

AWS MANAGEMENT CONSOLE



October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

L5.11

BILLING AND SECURITY

- My Billing Dashboard
  - Bills
  - Credits
  - Consolidated Billing
- My Security Credentials
  - Groups – collections of common permissions
    - More typically for a service “at large”
  - Users – Assigned to groups, have Access Key & Secret Key (AWS CLI), Signing Certificate (used for some APIs)
  - Roles – collect a group of security policies together
  - Policies – specific permissions to access services

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L10.12

## CLOUDWATCH

- Integration with many AWS services
- CloudTrail – enable detailed persisted logging to S3 for Lambda, other services
- Configure a Billing Alarm
- CloudWatch Events
- CloudWatch Rules and targets
- CloudWatch Metrics

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.13

## AWS EC2

- Elastic Compute Cloud
- Instance types: <https://ec2instances.info>
  - On demand instance – full price
  - Reserved instance – contract based
  - Spot instance – auction based, terminates with 2 minute warning
  - Dedicated/reserved host – reserved HW
  - Reserved host
  - Instance families:  
 General, compute-optimized, memory-optimized, GPU, etc.

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.14

## AWS EC2 - 2

- Internet connectivity
- Security Groups (firewall)
- Storage types
  - Instance storage - ephemeral storage
    - 1<sup>st</sup>, 2<sup>nd</sup>, and 3<sup>rd</sup> generation (e.g. c3/m3), and now c5d/m5d
  - EBS - Elastic block store
  - EFS - Elastic file system

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.15

## 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.16

## ELASTIC BLOCK STORE

- EBS cost model is different than instance storage (uses S3)
  - ~10¢ per GB/month
  - 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.17

## EBS VOLUME TYPES - 2

- **Metric: I/O Operations per Second (IOPS)**
- General Purpose 2 (GP2)
  - 3 IOPS per GB, Max 10,000 IOPS, 160MB/sec per volume
- Provisioned IOPS (IO1)
  - 32,000 IOPS, and 500 MB/sec throughput per volume
- 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 800 MB/sec throughput
  - 5 ¢ per GB/month

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.18

## ELASTIC FILE SYSTEM (EFS)

- Network file system (based on NFSv4 protocol)
- Shared file system for EC2 instances
- 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L10.19

## ELASTIC FILE SYSTEM (EFS) - 2

- Burstable Throughput Rates**
  - 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.20

## ELASTIC FILE SYSTEM (EFS) - 3

- Throughput Models**
- 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.21

## ELASTIC FILE SYSTEM (EFS) - 4

Performance Comparison, Amazon EFS and Amazon EBS

	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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.22

## 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L10.23

## EC2 VIRTUALIZATION - PARAVIRTUAL

- 1<sup>st</sup>, 2<sup>nd</sup>, 3<sup>rd</sup>, 4<sup>th</sup> generation → XEN-based**
- 5<sup>th</sup> 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 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L10.24

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 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

LS.25

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 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

LS.26

EVOLUTION OF AWS VIRTUALIZATION

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

AWS EC2 Virtualization Types

#	Tech	Type	With	VS	VS	VS	VS	VS	VS
1	VM	Fully Emulated		VS	VS	VS	VS	VS	VS
2	VM	Xen PV 3.0	PV drivers	P	P	P	P	VS	VS
3	VM	Xen HVM 3.0	PV drivers	VH	P	P	P	VS	VS
4	VM	Xen HVM 4.0.1	PVHVM drivers	VH	P	P	P	P	VS
5	VM	Xen AWS 2013	PVHVM + SR-IOV (net)	VH	VH	P	P	P	VS
6	VM	Xen AWS 2017	PVHVM + SR-IOV (net, stor)	VH	VH	VH	P	P	VS
7	VM	AWS Nitro 2017		VH	VH	VH	VH	VH	VS
8	HW	AWS Bare Metal 2017		H	H	H	H	H	H

VM: Virtual Machine, HW: Hardware, VS: Virt. in software, VH: Virt. in hardware, P: Paravirt, Not all combinations shown, SR-IOV (net): sgubleria driver, SR-IOV (stor): nvme driver.

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

LS.27

INSTANCE ACTIONS

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

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

LS.28

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 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
Institute of Technology, University of Washington - Tacoma

LS.29

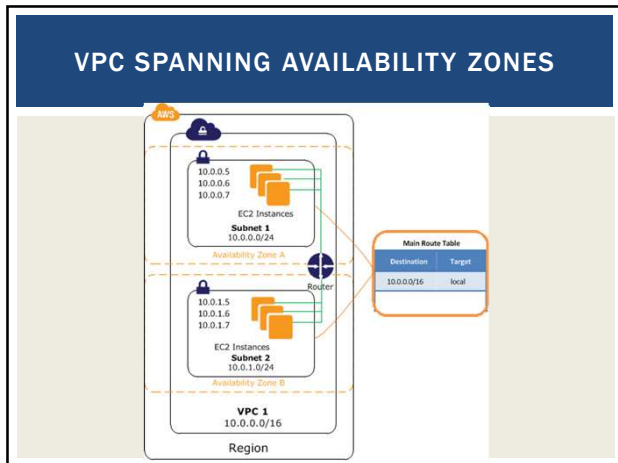
SIMPLE VPC

- Recommended when using Amazon EC2

October 29, 2018

TCSS562: Software Engineering for Cloud Computing [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

LS.30



### VIRTUAL PRIVATE CLOUD (VPC)

- Core components
  - VPCs
  - Subnets
  - Route Tables
  - Internet Gateways
  - DHCP Option Sets
  - Elastic IPs
  - NAT Gateways

October 29, 2018 TCS5562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma L10.32

### 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 29, 2018 TCS5562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma L10.33

### 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 29, 2018 TCS5562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma L10.34

### AWS CLI - 2

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

October 29, 2018 TCS5562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma L10.35

### 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 29, 2018 TCS5562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma L10.36

## ALTERNATIVE CLI

- `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`

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.37

## INSPECTING INSTANCE INFORMATION

- 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 (??)

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.38

## 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#icmipid=docs_iam_console#ami-tools-create-certificate`

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.39

## 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:
  1. Configure VM with software stack
  2. Burn new image for VM replication (**horizontal scaling**)
- Some folks may just install Docker. . .
- Create image script . . .

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.40

## CREATE A NEW INSTANCE STORE IMAGE SCRIPT

```
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/amiutils/cert-ec2.pem --no-inherit -r x86_64 -p $image -i
/etc/ec2/amiutils/cert-ec2.pem
cd /tmp
ec2-upload-bundle -b tc5562 -m $image.manifest.xml -a ${AWS_ACCESS_KEY} -s
${AWS_SECRET_KEY} --url http://s3.amazonaws.com --location US
ec2-register tc5562/$image.manifest.xml --region us-east-1 --kernel aki-
88aa75e1
```

October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 Institute of Technology, University of Washington - Tacoma

L5.41

## QUESTIONS



October 29, 2018

TCCS562: Software Engineering for Cloud Computing [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L10.42