

TCSS 562: SOFTWARE ENGINEERING FOR CLOUD COMPUTING

Cloud Computing Concepts and Models



Wes J. Lloyd
 Institute of Technology
 University of Washington - Tacoma

OBJECTIVES

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

March 30, 2017 TCSS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.2

ROLES

- **Cloud provider**
 - Organization that provides cloud-based resources
 - Responsible for fulfilling SLAs for cloud services
 - Some cloud providers "resell" IT resources from other cloud providers
 - Example: Heroku sells PaaS services running atop of Amazon EC2
- **Cloud consumers**
 - Cloud users that consume cloud services
- **Cloud service owner**
 - Both cloud providers and cloud consumers can own cloud services
 - A cloud service owner may use a cloud provider to provide a cloud service (e.g. Heroku)

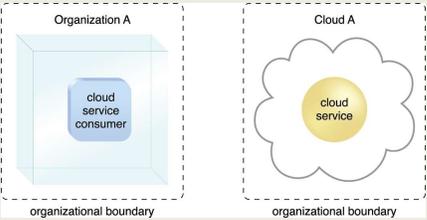
March 30, 2017 TCSS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.3

ROLES - 2

- **Cloud resource administrator**
 - Administrators provide and maintain cloud services
 - Both cloud providers and cloud consumers have administrators
- **Cloud auditor**
 - Third-party which conducts independent assessments of cloud environments to ensure security, privacy, and performance.
 - Provides unbiased assessments
- **Cloud brokers**
 - An intermediary between cloud consumers and cloud providers
 - Provides service aggregation
- **Cloud carriers**
 - Network and telecommunication providers which provide network connectivity between cloud consumers and providers

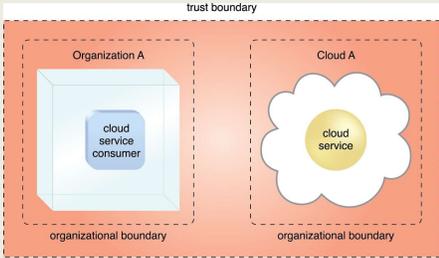
March 30, 2017 TCSS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.4

ORGANIZATION BOUNDARY



March 30, 2017 TCSS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.5

TRUST BOUNDARY



March 30, 2017 TCSS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.6

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

March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.7

ON-DEMAND USAGE

- The freedom to self-provision IT resources
- Generally with automated support
- Automated support requires no human involvement
- Automation through software services interface



March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.8

UBIQUITOUS ACCESS

- Cloud services are widely accessible
- Public cloud: internet accessible
- Private cloud: throughout segments of a company's intranet
- 24/7 availability

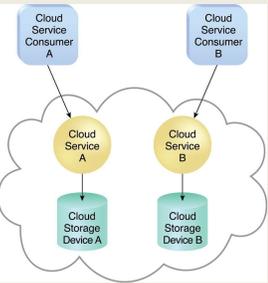
March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.9

MULTITENANCY

- Cloud providers pool resources together to share them with many users
- Serve multiple cloud service consumers
- IT resources can be dynamically assigned, reassigned based on demand
- Multitenancy can lead to performance variation

March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.10

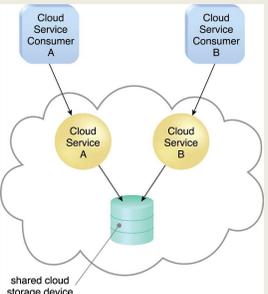
SINGLE TENANT MODEL



March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.11

MULTITENANT MODEL

- Resource is "multiplexed" and share amongst multiple users
- Goal is to increase utilization
- Often server resources are underutilized
- There are many "sunk costs" whether usage is 0% or 100%
- Cloud computing tries to maximize "sunk cost" investments



March 30, 2017
TCCS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
L2.12

MULTITENANT DATABASE

Isolated

Tenant A
Tenant B
Tenant C

Separate database

E1

Semi-shared

Tenant A
Tenant B
Tenant C

Shared database
Separate schema

E2

Shared

Tenant A
Tenant B
Tenant C

Shared database
Shared schema

E3

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.13

MULTITENANCY OF RESOURCES

Where is the multitenancy?

Traditional On Premise

Users
LAN
Application
Database
Server

Single Tenant (Hosted)

Customer
Customer
VPN
VPN
Application
Application
Database
Database
Server

Multi-Tenant

Customer
Customer
SSL
SSL
Application
Database
Server

Virtual Appliance

Customer
Customer
SSL
SSL
Application
Application
Database
Database
Virtualization
Virtualization
Server

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.14

ELASTICITY

- Automated ability of cloud to transparently scale resources
- Scaling based on runtime conditions or pre-determined by cloud consumer or cloud provider
- Threshold based scaling
 - CPU-utilization > threshold_A, Response_time > 100ms
 - Application agnostic vs. application specific thresholds
 - Why might an application agnostic threshold be non-ideal?
- Load prediction
 - Historical models
 - Real-time trends

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.15

PREDICTABLE DEMAND

Example:

Auto-Scaling Example: Netflix

Avg CPU
Instances

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.16

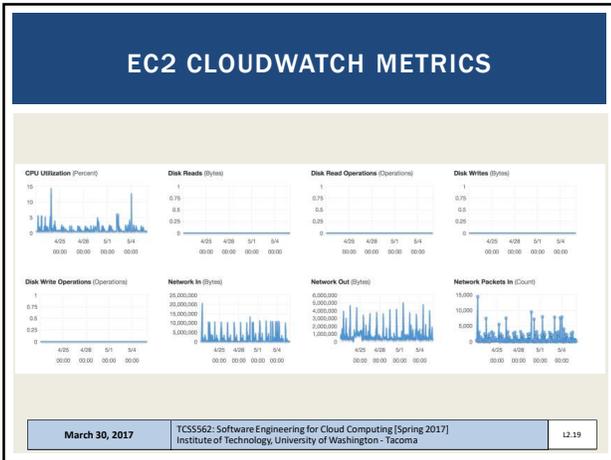
MEASURED USAGE

- Cloud platform tracks usage of IT resources
- For billing purposes
- Enables charging only for IT resources actually used
- Can be time-based (minute, hour, day)
- Can be throughput-based (MB, GB)
- Not all measurements are for billing
- Some measurements can support auto-scaling
- For example CPU utilization

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.17

EC2 CLOUDWATCH METRICS

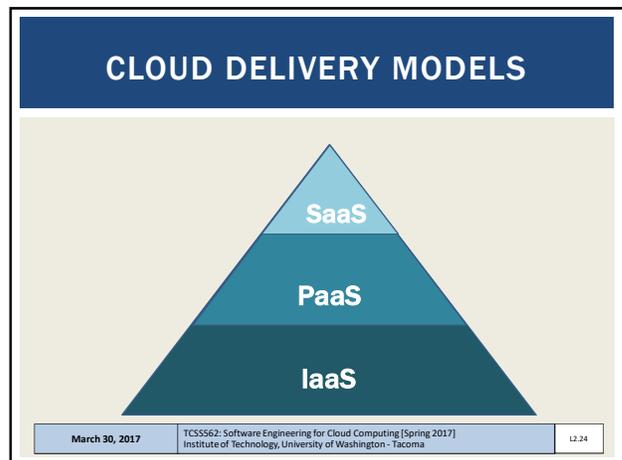
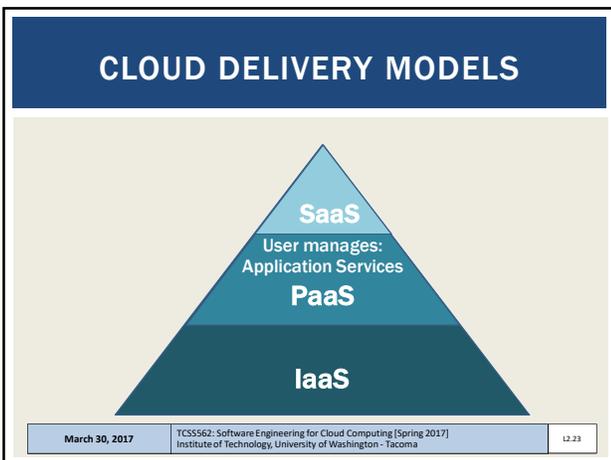
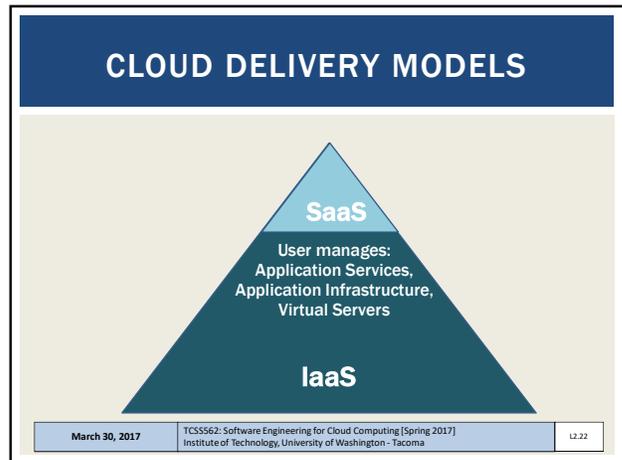
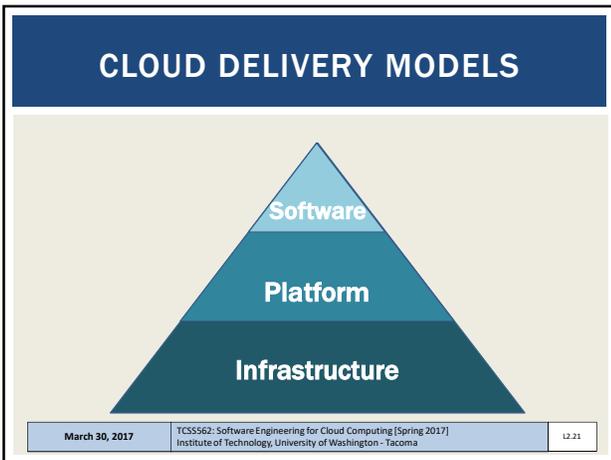
March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma L2.18



RESILIENCY

- Distributed redundancy across physical locations
- 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

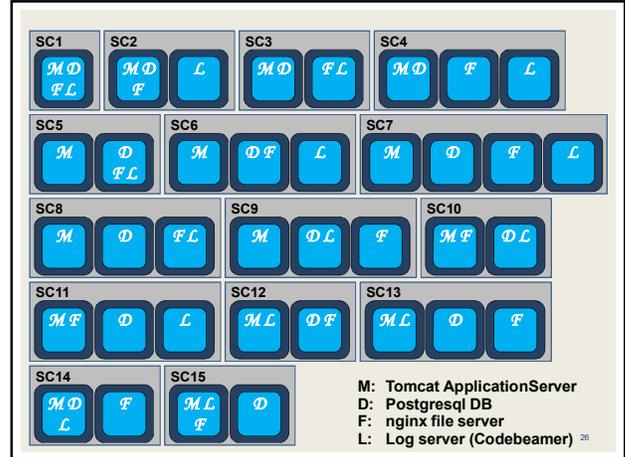
March 30, 2017 | TCCS562: Software Engineering for Cloud Computing [Spring 2017]
 Institute of Technology, University of Washington - Tacoma | L2.20



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

March 30, 2017
TCSS562: Software Engineering for Cloud Computing [Spring 2017]
Institute of Technology, University of Washington - Tacoma
12.25



Bell's Number:

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

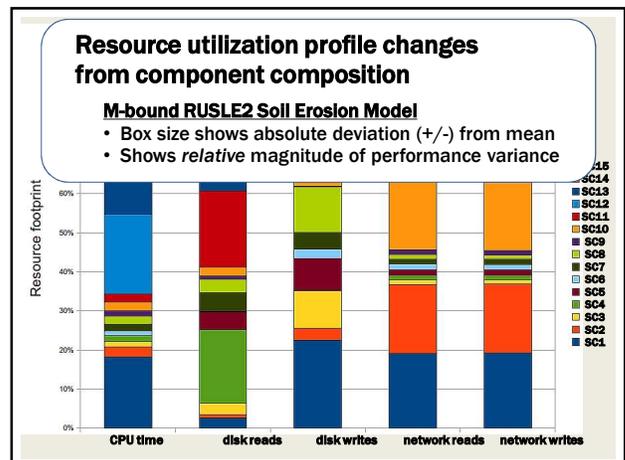
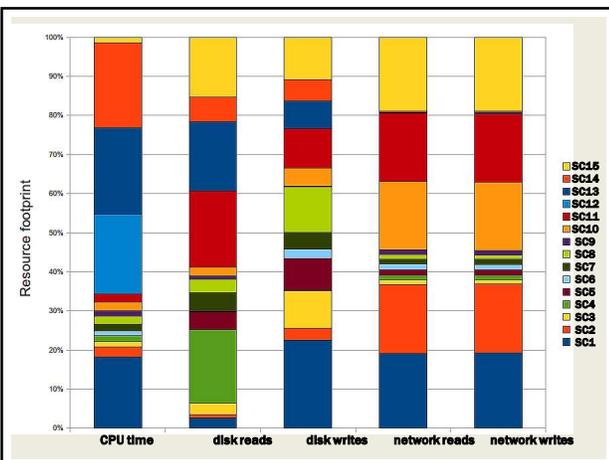
n	k
4	15
5	52
6	203
7	877
8	4,140
9	21,147
n	...

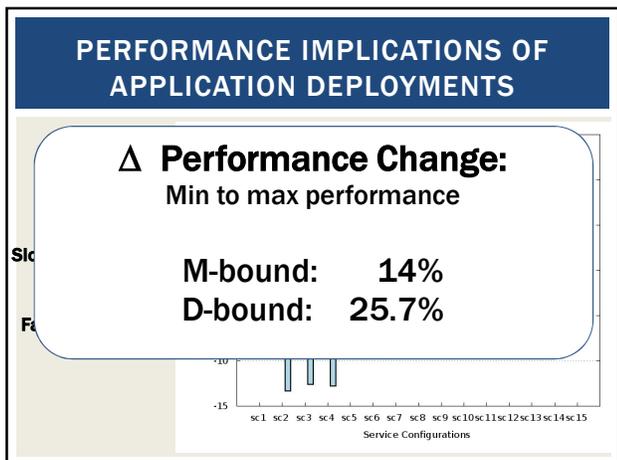
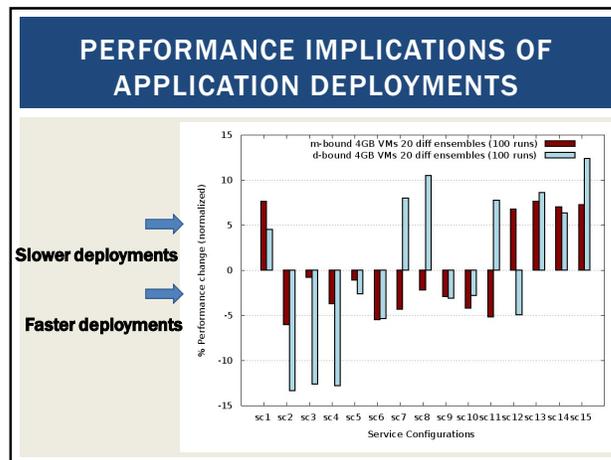
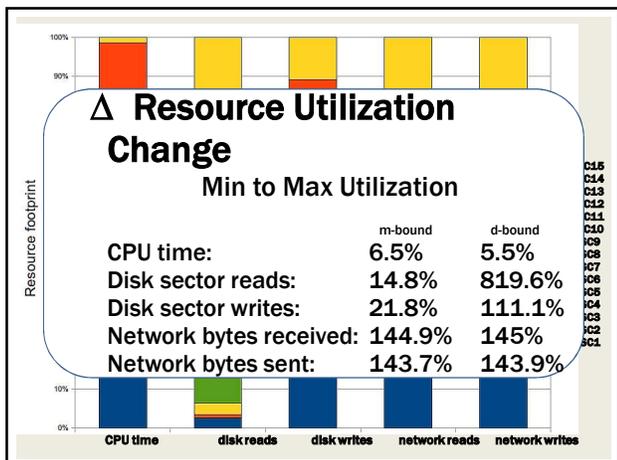
SC14: M, D, L
SC15: M, L, D
M: Tomcat ApplicationServer
D: Postgresql DB
F: nginx file server
L: Log server (Codebeamer) 27

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

SC14: M, D, L
SC15: M, L, D
M: Tomcat ApplicationServer
D: Postgresql DB
F: nginx file server
L: Log server (Codebeamer) 28





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

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017] Institute of Technology, University of Washington - Tacoma L2.34

AWS LAMBDA

How service requests influenced infrastructure:

Requests per batch	1	10	20	50	100
VMs/request Standard	1	1.72	2.94	5	5.56
VMs/request Oregon	1	2.78	3.12	5	5.56

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017] Institute of Technology, University of Washington - Tacoma L2.35

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

March 30, 2017 TCCS562: Software Engineering for Cloud Computing [Spring 2017] Institute of Technology, University of Washington - Tacoma L2.36

INTRODUCTIONS

- **SPEAK SLOWLY**
- Name
- What year of your program (quarter) are you in?
 - E.g. senior, 2nd quarter of MS, etc.
- Where are you from?
- What are your computer science and/or cloud/project interests?

March 30, 2017	TCSS562: Software Engineering for Cloud Computing [Spring 2017] Institute of Technology, University of Washington - Tacoma	L2.37
----------------	---	-------

QUESTIONS



March 30, 2017	TCSS562: Software Engineering for Cloud Computing [Spring 2017] Institute of Technology, University of Washington - Tacoma	L2.38
----------------	---	-------