

TCSS 562: SOFTWARE ENGINEERING FOR CLOUD COMPUTING

Containerization, Kubernetes



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

MW 5:50-7:50 PM

1

OBJECTIVES – 11/23

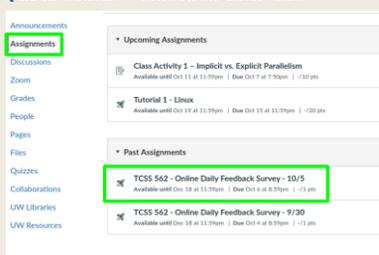
- **Questions from 11/18**
- Quiz 2 – to be posted ~ Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 – Intro to Docker/Containerization
- **Group Presentation Overview:**
 Cloud Technology or Research Paper for 11/30 – 12/9
- Term Project Check-in – due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- **2nd hour:**
- Introduction to Kubernetes cont'd
- Tutorial questions / Team planning

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington – Tacoma LIS.2

2

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas – Take After Each Class
- Extra Credit for completing



November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington – Tacoma LIS.3

3

TCSS 562 - Online Daily Feedback Survey - 10/5

Started Oct 7 at 1:13pm

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

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington – Tacoma LIS.4

4

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (27 respondents):
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average – 6.48 (↑ - previous 6.30)**
- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average – 5.37 (↑ - previous 5.35)**

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington – Tacoma LIS.5

5

FEEDBACK FROM 11/18

- **Is there a major difference between the performance of containerization tools (e.g. Docker vs Podman)? Has there been any recent benchmarks that compares performance of different container tools and/or security?**
 - I'm aware of and have newer publications from 2015 and beyond. Some papers have recently focused more on MicroVMs
 - Am happy to share these for a cloud paper presentation
 - An updated study that investigates Docker/Podman, gVisor, Kata Containers, Nabla, and others is needed
 - Possible capstone / thesis topic
- To gain a sense regarding how diverse the various tools have become see this blog:
"Welcome to the Container Jungle: Docker vs. containerd vs. Nabla vs. Kata vs. Firecracker and more!"
- <https://www.inovex.de/de/blog/containers-docker-containerd-nabla-kata-firecracker/>

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington – Tacoma LIS.6

6

UPCOMING TUTORIALS

- Tutorial 7 – **[POSTED]** Introduction to Docker Containerization
- **Extra credit tutorials – submit by Dec 17 @ 11:59p**
- Tutorial 8 – Introduction to FaaS IV: Step Functions and SQS
- Tutorial 9 – Asynchronous Function Profiling with SAAF
- **Ungraded tutorials:**
- Tutorial 10 – Automating Experiments with SAAF & FaaS Runner
- There may be 1 other optional tutorial posted...

November 23, 2021
TCSS562: Software Engineering for Cloud Computing (Fall 2021)
School of Engineering and Technology, University of Washington - Tacoma
L15.7

7

OBJECTIVES – 11/23

- Questions from 11/18
- **Quiz 2 – to be posted ~ Dec 6**
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 – Intro to Docker/Containerization
- **Group Presentation Overview:**
Cloud Technology or Research Paper for 11/30 – 12/9
- Term Project Check-in – due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- **2nd hour:**
- Introduction to Kubernetes cont'd
- Tutorial questions / Team planning

November 23, 2021
TCSS562: Software Engineering for Cloud Computing (Fall 2021)
School of Engineering and Technology, University of Washington - Tacoma
L15.8

8

OBJECTIVES – 11/23

- Questions from 11/18
- Quiz 2 – to be posted ~ Dec 6
- **Tutorial 6 - Intro to FaaS III - Serverless Databases**
- Tutorial 7 – Intro to Docker/Containerization
- **Group Presentation Overview:**
Cloud Technology or Research Paper for 11/30 – 12/9
- Term Project Check-in – due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- **2nd hour:**
- Introduction to Kubernetes cont'd
- Tutorial questions / Team planning

November 23, 2021
TCSS562: Software Engineering for Cloud Computing (Fall 2021)
School of Engineering and Technology, University of Washington - Tacoma
L15.9

9

OBJECTIVES – 11/23

- Questions from 11/18
- Quiz 2 – to be posted ~ Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- **Tutorial 7 – Intro to Docker/Containerization**
- **Group Presentation Overview:**
Cloud Technology or Research Paper for 11/30 – 12/9
- Term Project Check-in – due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- **2nd hour:**
- Introduction to Kubernetes cont'd
- Tutorial questions / Team planning

November 23, 2021
TCSS562: Software Engineering for Cloud Computing (Fall 2021)
School of Engineering and Technology, University of Washington - Tacoma
L15.10

10

TUTORIAL COVERAGE

- **Tutorial Concepts:**
- Docker installation
- Working with docker files
- Publishing images to Docker Hub
- **Docker CLI:**
 - Docker run – create a container
 - Docker ps – list containers
 - Docker exec –it – run a process in an existing container
 - Docker stop –stop container
 - Docker image
- Using cgroups to inspect container resource utilization metrics
- Container resources quotas: memory, CPU
- Testing CPU and memory isolation of co-located containers

November 23, 2021
TCSS562: Software Engineering for Cloud Computing (Fall 2021)
School of Engineering and Technology, University of Washington - Tacoma
L15.11

11

Docker CLI

```

attach      Attach local standard input, output, and error streams to a running container
build       Build an image from a Dockerfile
commit     Create a new image from a container's changes
cp         Copy files/folders between a container and the local filesystem
create     Create a new container
deploy     Deploy a new stack or update an existing stack
diff       Inspect changes to files or directories on a container's filesystem
events     Get real time events from the server
exec       Run a command in a running container
export     Export a container's filesystem as a tar archive
history    Show the history of an image
images     List images
import     Import the contents from a tarball to create a filesystem image
info       Display system-wide information
inspect    Return low-level information on Docker objects
kill       Kill one or more running containers
load      Load an image from a tar archive or STDIN
login      Log in to a Docker registry
logout    Log out from a Docker registry
logs      Fetch the logs of a container
pause     Pause all processes within one or more containers
port      List port mappings or a specific mapping for the container
ps        List containers
pull      Pull an image or a repository from a registry
push      Push an image or a repository to a registry
rename    Rename a container
restart   Restart one or more containers
rm        Remove one or more containers
rmi       Remove one or more images
run       Run a command in a new container
save      Save one or more images to a tar archive (streamed to STDOUT by default)
search    Search the Docker Hub for images
start     Start one or more stopped containers
stats     Display a live stream of container(s) resource usage statistics
stop     Stop one or more running containers
tag       Create a tag TARGET_IMAGE that refers to SOURCE_IMAGE
top       Display the running processes of a container
unpause  Unpause all processes within one or more containers
update    Update configuration of one or more containers
version   Show the Docker version information
wait     Block until one or more containers stop, then print their exit codes
    
```

12

TUTORIAL 7

- Linux performance benchmarks
- stress-ng**
 - 100s of CPU, memory, disk, network stress tests
 - Designed to generate a load on various parts of a system
- Sysbench**
 - Common system benchmark included with Linux
 - Used in tutorial 7 for memory stress test

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.13

13

OBJECTIVES - 11/23

- Questions from 11/18
- Quiz 2 - to be posted ~ Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 - Intro to Docker/Containerization
- Group Presentation Overview:
Cloud Technology or Research Paper for 11/30 - 12/9**
- Term Project Check-in - due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- 2nd hour:**
 - Introduction to Kubernetes cont'd
 - Tutorial questions / Team planning

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.14

14

GROUP PRESENTATION

- TWO OPTIONS:**
- Cloud technology presentation**
- Cloud research paper presentation**
 - Recent & suggested papers will be posted at:
<http://faculty.washington.edu/wlloyd/courses/tcss562/papers/>
- Submit presentation type and topics (paper or technology) with desired dates of presentation via Canvas by:
TODAY: Tuesday November 23rd @ 11:59pm**
- Presentation dates:**
 - Tuesday November 30, Thursday December 2
 - Tuesday December 7, Thursday December 9

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.15

15

OBJECTIVES - 11/23

- Questions from 11/18
- Quiz 2 - to be posted ~ Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 - Intro to Docker/Containerization
- Group Presentation Overview:
Cloud Technology or Research Paper for 11/30 - 12/9**
- Term Project Check-in - due Thur 12/2 @ 11:59p**
- Introduction to Containerization cont'd
- Introduction to Kubernetes
- 2nd hour:**
 - Introduction to Kubernetes cont'd
 - Tutorial questions / Team planning

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.16

16

OBJECTIVES - 11/23

- Questions from 11/18
- Quiz 2 - to be posted ~ Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 - Intro to Docker/Containerization
- Group Presentation Overview:
Cloud Technology or Research Paper for 11/30 - 12/9**
- Term Project Check-in - due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd**
- Introduction to Kubernetes
- 2nd hour:**
 - Introduction to Kubernetes cont'd
 - Tutorial questions / Team planning

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.17

17

CONTAINERIZATION



November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma | LIS.18

18

OTHER DOCKER TOOLS

- **Docker Machine:** automatically provision and manage sets of docker hosts to form a cluster
- **Docker Swarm:** Clusters multiple docker hosts together to manage as a cluster.
- **Docker Compose:** Config file (YAML) for multi-container application; Describes how to deploy and configure multiple containers

```

            graph TD
            DE[Docker Engine] --> C[containerd]
            C --> CS1[containerd-shim]
            C --> CS2[containerd-shim]
            C --> CS3[...]
            CS1 --> R1[runC]
            CS2 --> R2[runC]
            CS3 --> R3[...]
            
```

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.19

19

CONTAINER ORCHESTRATION FRAMEWORKS

- Framework(s) to deploy multiple containers
- Provide container clusters using cloud VMs
- Similar to “private clusters”
- Reduce VM idle CPU time in public clouds
- Better leverage “sunk cost” resources
- Compact multiple apps onto shared public cloud infrastructure
- Generate to cost savings
- Reduce vendor lock-in

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.20

20

KEY ORCHESTRATION FEATURES

- Management of container hosts
- Launching set of containers
- Rescheduling failed containers
- Linking containers to support workflows
- Providing connectivity to clients outside the container cluster
- Firewall: control network/port accessibility
- Dynamic scaling of containers: horizontal scaling
 - Scale in/out, add/remove containers
- Load balancing over groups of containers
- Rolling upgrades of containers for application

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.21

21

CONTAINER ORCHESTRATION FRAMEWORKS - 2

- Docker swarm
- Apache mesos/marathon
- Kubernetes
 - Many public cloud providers moving to offer Kubernetes-as-a-service
- Amazon elastic container service (ECS)
- Apache aurora
- Container-as-a-Service
 - Serverless containers without managing clusters
 - Azure Container Instances, AWS Fargate...

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.22

22

OBJECTIVES - 11/23

- Questions from 11/18
- Quiz 2 - to be posted - Dec 6
- Tutorial 6 - Intro to FaaS III - Serverless Databases
- Tutorial 7 - Intro to Docker/Containerization
- **Group Presentation Overview:** Cloud Technology or Research Paper for 11/30 - 12/9
- Term Project Check-in - due Thur 12/2 @ 11:59p
- Introduction to Containerization cont'd
- **Introduction to Kubernetes**
- **2nd hour:**
 - Introduction to Kubernetes cont'd
 - Tutorial questions / Team planning

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.23

23

OBJECTIVES - 11/23

- Questions from 11/18
- Quiz 2 - to be posted ~ Dec 6
- No Office Hours 11/25
- Class on 11/25: Office hours, and finish any remaining lecture from today
- Introduction to Containerization cont'd
- **Tutorial 7**
- **2nd hour:**
 - Introduction to Kubernetes
 - Tutorial questions
 - Team planning

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma
L15.24

24



KUBERNETES

from: "The Kubernetes Book", Nigel Poulton and Pushkar Joglekar, Version 7.0, September 2020

L15.25

25

KUBERNETES

- Name is from the Greek word meaning Helmsman
 - The person who steers a seafaring ship
 - The logo reinforces this theme
- Kubernetes is also sometimes called K8s
- Kubernetes is an application orchestrator



- Most common use case is to containerize cloud-native microservices applications
- What is an orchestrator?
 - System that deploys and manages applications

November 23, 2021 TCCS562: Software Engineering for Cloud Computing [Fall 2021] School of Engineering and Technology, University of Washington - Tacoma L15.26

26

KUBERNETES - 2

Why does Google want to give Kubernetes away for free?

- Initially developed by Google
- **Goal:** make it easier for potential customers to use Google Cloud
- Kubernetes leverages knowledge gained from two internal container management systems developed at Google
 - Borg and Omega
- Google donated Kubernetes to the Cloud Native Computing Foundation in 2014 as an open-source project
- Kubernetes is written in Go (Golang)
- Kubernetes is available under the Apache 2.0 license
- Releases were previously maintained for only 8 months!
- Starting w/ v 1.19 (released Aug 2020) support is 1 year

November 23, 2021 TCCS562: Software Engineering for Cloud Computing [Fall 2021] School of Engineering and Technology, University of Washington - Tacoma L15.27

27

GOALS OF KUBERNETES

1. Deploy your application
2. Scale it up and down dynamically according to demand
3. Self-heal it when things break
4. Perform zero-downtime rolling updates and rollbacks

- These features represent automatic infrastructure management
- Containerized applications run in container(s)
- Compared to VMs, containers are thought of as being:
 - Faster
 - More light-weight
 - More suited to rapidly evolving software requirements

November 23, 2021 TCCS562: Software Engineering for Cloud Computing [Fall 2021] School of Engineering and Technology, University of Washington - Tacoma L15.28

28

CLOUD NATIVE APPLICATIONS

- Applications designed to meet modern software requirements including:
 - **Auto-scaling:** resources to meet demand
 - **Self-healing:** required for high availability (HA) and fault tolerance
 - **Rolling software updates:** with no application downtime for DevOPS
 - **Portability:** can run anywhere there's a Kubernetes cluster

November 23, 2021 TCCS562: Software Engineering for Cloud Computing [Fall 2021] School of Engineering and Technology, University of Washington - Tacoma L15.29

29

WHAT IS A MICROSERVICES APP?

- Application consisting of many specialized parts that communicate and form a meaningful application
- Example components of a microservice eCommerce app:

Web front-end	Catalog service
Shopping cart	Authentication service
Logging service	Persistent data store
- **KEY IDEAS:**
 - Each microservice can be coded/maintained by different team
 - Each has its own release cadence
 - Each is deployed/scaled separately
 - Can patch & scale the log service w/o impacting others

November 23, 2021 TCCS562: Software Engineering for Cloud Computing [Fall 2021] School of Engineering and Technology, University of Washington - Tacoma L15.30

30

KUBERNETES - 3

- Provides "an operating system for the cloud"
- Offers the de-facto standard platform for deploying and managing cloud-native applications
- OS: abstracts physical server, schedules processes
- Kubernetes: **abstracts the cloud**, schedules microservices
- Kubernetes abstracts differences between private and public clouds
- Enable cloud-native applications to be cloud agnostic
 - i.e. they don't care WHAT cloud they run on
 - Enables fluid application migration between clouds
- Kubernetes provides rich set of tools/APIs to introspect (observe and examine) your apps

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.31

31

KUBERNETES - 4

- Features:
 - A "control plane" – brain of the cluster
 - Implements autoscaling, rolling updates w/o downtime, self-healing
 - A "bunch of nodes" – workers (muscle) of the cluster
- Provides orchestration
- The process of organizing everything into a useful application
- And also keeping it running smoothly

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.32

32

KUBERNETES - CLUSTER MANAGEMENT

- Master node(s) manage the cluster by:
 - Making scheduling decisions
 - Performing monitoring
 - Implementing changes
 - Responding to events
- Masters implement the control plane of a Kubernetes cluster
- Recipe for deploying to Kubernetes:
 - Write app as independent microservices in preferred language
 - Package each microservice in a container
 - Create a manifest to encapsulate the definition of a Pod
 - Deploy Pods to the cluster w/ a higher-level controller such as "Deployments" or "DaemonSets"

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.33

33

DECLARATIVE SERVICE APPROACH

- **Imperative definition:** sets of commands and operations
 - Example: BASH script, Dockerfile
- **Declarative definition:** specification of a service's properties
 - What level of service it should sustain, etc.
 - Example: Kubernetes YAML files
- Kubernetes manages resources **declaratively**
- How apps are deployed and run are defined with YAML files
- YAML files are POSTed to Kubernetes endpoints
- Kubernetes deploys and manages applications based on declarative service requirements
- If something isn't as it should be: *Kubernetes automatically tries to fix it*

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.34

34

KUBERNETES MASTERS

- Provide system services to host the control plane
- Simplest clusters use only 1 master – no replication
 - Suitable for lab and dev/test environments
- Production environments: masters are replicated ~3-5x
 - Provides fault tolerance and high availability (HA)
 - Cloud-based managed Kubernetes services offer HA deployments

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.35

35

MASTER SERVICES

- **API Server**
- Cluster store
- Controller Manager
- Scheduler
- Cloud controller

November 23, 2021 TCSS562: Software Engineering for Cloud Computing [Fall 2021]
 School of Engineering and Technology, University of Washington - Tacoma LIS.36

36

API SERVER

- Can run on 1-node for lab, test/dev environments
- Default port is 443
- Exposes a RESTful API where YAML configuration files are POST(ed) to
- YAML files (manifests) describe desired state of an application
 - Which container image(s) to use
 - Which ports to expose
 - How many POD replicas to run

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.37

37

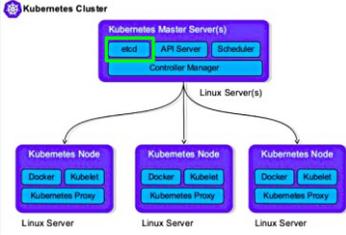
WE WILL RETURN AT ~6:10PM



38

MASTER SERVICES

- API Server
- Cluster store**
- Controller Manager
- Scheduler
- Cloud controller



November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.39

39

CLUSTER STORE

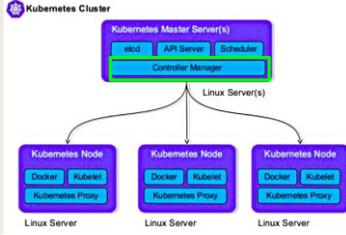
- Used to persist Kubernetes cluster state
- Persistently stores entire configuration and state of the cluster
- Currently implemented with **etcd**
 - Popular distributed key/value store (db) supporting replication
 - HA deployments may use ~3-5 replicas
 - Is the authority on true state of the cluster
- etcd prefers consistency over availability
- etcd failure: apps continue to run, nothing can be reconfigured
- Consistency of writes is vital
- Employs **RAFT consensus protocol** to negotiate which replica has correct view of the system in the event of replica failure

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.40

40

MASTER SERVICES

- API Server
- Cluster store
- Controller Manager**
- Scheduler
- Cloud controller



November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.41

41

CONTROLLER MANAGER

- Provides a "controller" of the controllers
 - Implements background control loops to monitor cluster and respond to events
 - Control loops include: node controller, endpoints controller, replicaset controller, etc...
- GOAL: ensure cluster current state matches desired state**
- Control Loop Logic:**
 - Obtain desired state (defined in manifest YAMLs)
 - Observe the current state
 - Determine differences
 - Reconcile differences
- Controllers are specialized to manage a specific resource type
 - They are not aware/concerned with of other parts of the system

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.42

42

MASTER SERVICES

- API Server
- Cluster store
- Controller Manager
- Scheduler
- Cloud controller

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.43

43

TASK SCHEDULER

- Scheduler's job is to identify the best node to run a task
 - Scheduler does not actually run tasks itself
- Assigns work tasks to appropriate healthy nodes
- Implements complex logic to filter out nodes incapable of running specified task(s)
- Capable nodes are ranked
- Node with highest ranking is selected to run the task

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.44

44

ENFORCING SCHEDULING PREDICATES

- Scheduler performs predicate (property) checks to verify how/where to run tasks
 - Is a node tainted?
 - Does task have affinity (deploy together), anti-affinity (separation) requirements?
 - Is a required network port available on the node?
 - Does node have sufficient free resources?
- Nodes incapable of running the task are eliminated as candidate hosts

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.45

45

RANKING NODES

- Remaining nodes are ranked based on for example:
 1. Does the node have the required images?
 - Cached images will lead to faster deployment time
 2. How much free capacity (CPU, memory) does the node have?
 3. How many tasks is the node already running?
- Each criterion is worth points
- **Node with most points is selected**
- If there is no suitable node, task is not scheduled, but marked as pending
- **PROBLEM:** *There is no one-sized fits all solution to selecting the best node. How weights are assigned to conditions may not reflect what is best for the task*

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.46

46

MASTER SERVICES

- API Server
- Cluster store
- Controller Manager
- Scheduler
- Cloud controller

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.47

47

CLOUD CONTROLLER MANAGER

- Abstracts and manages integration with specific cloud(s)
- Manages vendor specific cloud infrastructure to provide instances (VMs), load balancing, storage, etc.
- Support for AWS, Azure, GCP, Digital Ocean, IBM, etc.

November 23, 2021 | TCSS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | LIS.48

48

MASTER SERVICES

- API Server
- Cluster store
- Controller Manager
- Scheduler
- Cloud controller

The diagram shows a central 'Kubernetes Cluster' box containing 'Kubernetes Master Server(s)' with sub-components 'etcd', 'API Server', 'Scheduler', and 'Controller Manager'. Below this is a 'Linux Server(s)' box. Three arrows point from the Master Server(s) to three separate 'Kubernetes Node' boxes, each on a 'Linux Server'. Each node contains 'Kubernetes Proxy', 'Kubelet', and 'Docker'.

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.49

49

WORKER NODES

- Nodes perform tasks (i.e. host containers & services)
- Three primary functions:
 1. Wait for the scheduler to assign work
 2. Execute work (host containers, etc.)
 3. Report back state information, etc.
- Nodes are considerably simpler than masters

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.50

50

WORKER NODES

- Kubelet
- Container runtime (Docker, etc.)
- Kubernetes Proxy

The diagram is identical to slide 49, showing the Master Services and three Worker Nodes. The 'Kubelet' component within each 'Kubernetes Node' box is highlighted with a green border.

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.51

51

KUBELET

- Main Kubernetes agent
- Runs on every node
- Adding a new node installs the kubelet onto the node
- Kubelet registers the node with the cluster
- Monitors API server for new work assignments
- Maintains reporting back to control plane
- When a node can't run a task, kubelet is NOT responsible for finding an alternate node

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.52

52

WORKER NODES

- Kubelet
- Container runtime (Docker, etc.)
- Kubernetes Proxy

The diagram is identical to slide 51, showing the Master Services and three Worker Nodes. The 'Container runtime (Docker, etc.)' component within each 'Kubernetes Node' box is highlighted with a green border.

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.53

53

CONTAINER RUNTIME(S)

- Each node requires a container runtime to run containers
- Early versions had custom support for a limited number of container types, e.g. Docker
- Kubernetes now provides a standard Container Runtime Interface (CRI)
- CRI exposes a clean interface for 3rd party container runtimes to plug-in to
- Popular container runtimes: Docker, containerd, Kata

November 23, 2021 | TCS562: Software Engineering for Cloud Computing (Fall 2021) | School of Engineering and Technology, University of Washington - Tacoma | L15.54

54

WORKER NODES

- Kubelet
- Container runtime (*Docker, etc.*)
- **Kubernetes Proxy**

Kubernetes Cluster

Kubernetes Master Server(s)

etcd | API Server | Scheduler | Controller Manager

Linux Server(s)

Kubernetes Node | Docker | Kubelet | Kubernetes Proxy

Linux Server | Linux Server | Linux Server

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.55

55

KUBE-PROXY

- Runs on every node in the cluster
- Responsible for managing the cluster's networking
- Ensures each node obtains a unique IP address
- Implemented local IPTABLES and IPVS rules to route and load-balance traffic
- IPTABLES (ipv4) – enables configuration of IP packet filtering rules of the Linux kernel firewall
- IPVS – IP Virtual Server: provides transport-layer (layer 4) load balancing as part of the Linux kernel; Configured using ipvsadm tool in Linux

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.56

56

CORE KUBERNETES COMPONENTS

- **Kubernetes DNS**
- Pods
- Services

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.57

57

KUBERNETES DNS

- Every Kubernetes cluster has an internal DNS service
- Accessed with a static IP
- Hard-coded so that every container can find it
- Every service is registered with the DNS so that all components can find every Service on the cluster by **NAME**
- Is based on CoreDNS (<https://coredns.io>)

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.58

58

CORE KUBERNETES COMPONENTS

- Kubernetes DNS
- **Pods**
- Services

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.59

59

PODS

- Pod – atomic unit of deployment & scheduling in Kubernetes
- A Kubernetes Pod is defined to run a containerized application
- Kubernetes manages Pods, not individual containers
- Cannot run a container directly on Kubernetes
- All containers run through Pods
- Pod comes from "pod of whales"
- Docker logo shows a whale with containers stacked on top
- Whale represents the Docker engine that runs on a single host
- Pods encapsulate the definition of a single microservice for hosting purposes
- Pods can have a single container, or multiple containers if the service requires more than one

November 23, 2021 | TCS562: Software Engineering for Cloud Computing [Fall 2021] | School of Engineering and Technology, University of Washington - Tacoma | L15.60

60

PODS - 2

- Examples of multi-container Pods:
 - Service meshes
 - Web containers with a helper container that pulls latest content
 - Containers with a tightly coupled log scraper or profiler
- YAML manifest files are used to provide a declarative description for how to run and manage a Pod
- To run a pod, POST a YAML to the API Server: "kubectl run <NAME>" where NAME is the service
- A Pod runs on a single node (host)
- Pods share:
 - Interprocess communication (IPC) namespace
 - Memory, Volumes, Network stack

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.61

61

PODS - 3

- Pods provide a "fenced" environment to run containers
- Provide a "sandbox"
- Only tightly coupled containers are deployed with a single pod
- Best practice: decouple individual containers to separate pods
 - What is the best container composition into pods? (1:1, 1:many)
- **Scaling**
 - Pods are the unit of scaling
 - Add and remove pods to scale up/down
 - Do not add containers to a pod, add pod instances
 - Pod instances can be scheduled on the same or different host
- **Atomic Operation**
 - Pods are either fully up and running their service (i.e. port open/exposed), or pods are down / offline

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.62

62

PODS - 4

- **Pod Lifecycle**
 - An application should not be tightly bound or dependent on a specific Pod instance
 - Pods are designed to fail and be replaced
 - Use of **service objects** in Kubernetes help decouple pods to offer resiliency upon failure
- **Deployments**
 - Higher level controllers often used to deploy pods
 - Controllers implement a controller and watch loop:
 - "Deployments" – offer scalability & rolling updates
 - "DaemonSets" – run instance of service on every cluster node
 - "StatefulSets" – used for stateful components
 - "CronJobs" – for short lived tasks that need to run at specified times

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.63

63

CORE KUBERNETES COMPONENTS

- Kubernetes DNS
- Pods
- Services

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.64

64

KUBERNETES "SERVICES"

- Pods managed with "Deployments" or "DaemonSets" controllers are automatically replaced when they die
 - This provides resiliency for the application
- **KEY IDEA:** Pods are unreliable
- **Services** provide reliability by acting as a "GATEWAY" to pods that implement the services
 - They underlying pods can change over time
 - The services endpoints remain and are always available
- Service objects provide an abstraction layer w/ a reliable name and load balancing of requests to a set of pods

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.65

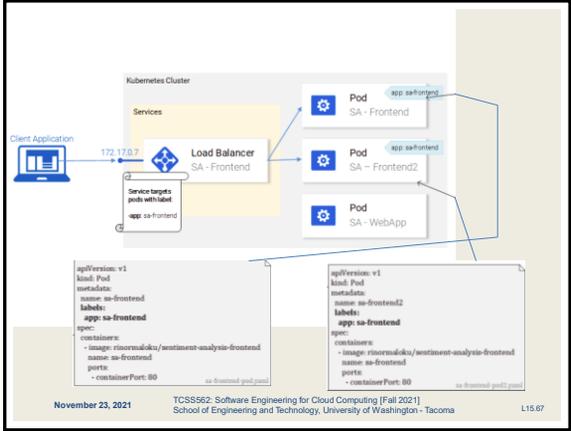
65

SERVICES

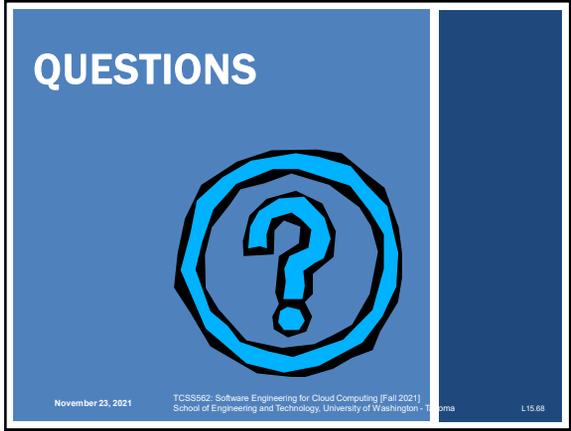
- Provide reliable front-end with:
 - Stable DNS name
 - IP Address
 - Port
- Services do not possess application intelligence
- No support for application-layer host and path routing
- Services have a "label selector" which is a set of labels
- Requests/traffic is only sent to Pods with matching labels
- Services only send traffic to healthy Pods
- **KEY IDEA:** Services bring stable IP addresses and DNS names to unstable Pods

November 23, 2021
TCSS562: Software Engineering for Cloud Computing [Fall 2021]
School of Engineering and Technology, University of Washington - Tacoma
L15.66

66



67



68