

AWS for Librarians 3: Cloud servers

vanderbi.it/learnaws

Steve Baskauf



Creating other kinds of serverless applications


- The output of one service can trigger a second service (e.g. Textract can extract text from an image, then Translate can translate that text to another language).
- **CRON** jobs can be used to trigger monitoring Lambdas that can pull data from an API and carry out some action
 - Example rainfall email, weather visualization
 - Plan to use this to pull API data (social media, etc.) into GitHub for the library dashboard.
- **Step functions** are a way to control serverless flow

Getting started with Docker

- Many servers are installed using Docker images
- the Community Edition is free and installation is straightforward (except when it doesn't work)
- Installation notes at
<https://heardlibrary.github.io/digital-scholarship/host/docker/>
- I had problems installing it on a Windows machine
- Docker Desktop runs in the background
- Docker commands are issued from the console (Terminal or Command Prompt)

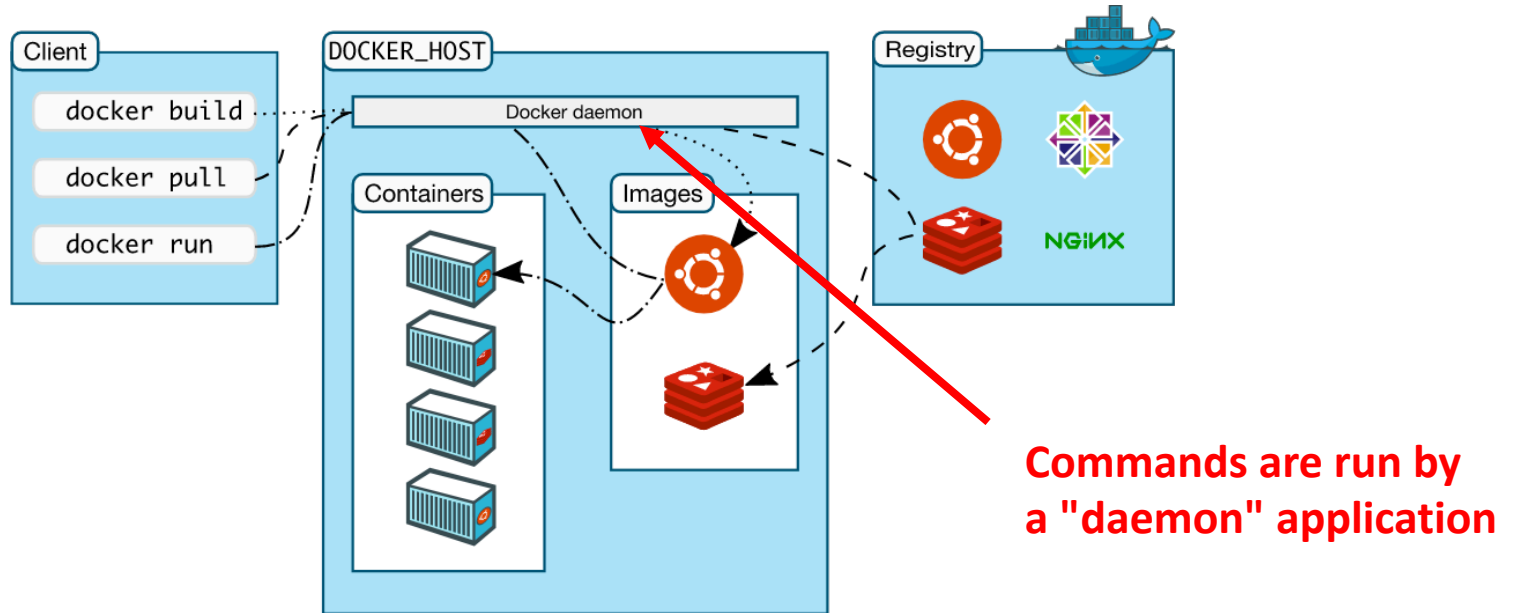
What is Docker?

- See <https://docs.docker.com/engine/docker-overview/>




- **Containers** are self-contained systems that include everything needed to run an application.
- **Images** are frozen, read-only versions of containers that can be used to spawn new containers.
- **Docker Hub** is a source of public images ("registry")

Docker commands



- **pull** pulls an image from Docker Hub to the environment
- **run** generates a container from an image and starts it running. It also associates a name with the container.
- **container stop** stops a running container
- **container restart** restarts an existing container
- **container/image ls** lists containers/images
- **container/image rm** deletes a container/image

Docker running on local computer



- Communication within the local computer through **localhost**: IP address
- The CLI talks to the daemon through a TCP/IP port
- Users talk to the container through some designated port.

A hand-drawn diagram illustrating a development environment architecture. On the left, a **client** (represented by a person) interacts with a **command line interface (CLI)** and a **docker-compose** box. The client is labeled **client** and has a red arrow pointing to it from the text **ker machine**. The client is connected via **Tcp/IP** to a **server** environment. The server is labeled **aws-sandbox (server) environment** and contains a **Docker daemon** (represented by a person) and an **nginx** box. The Docker daemon is connected to the **nginx** box. The server is also connected to a **port 8000** (indicated by a red arrow) and a **HTTP** connection. The server is connected to a **virtualbox** box, which is labeled **eval \$(docker-machine env aws-sandbox)**. The **virtualbox** box contains two **containers** (represented by small boxes). A red arrow points from the **virtualbox** box to the text **environments** on the right. The **client** is also connected to a **service link port** and a **digocan-app** box. The **client** is also connected to a **docker-compose** box. The **client** is also connected to a **service link port** and a **digocan-app** box. The **client** is also connected to a **service link port** and a **digocan-app** box.

- **Docker Machine** coordinates communication with multiple environments (local and remote)
- **Docker Compose** coordinates multiple containers in an environment (not necessary for single containers).

EC2 (Elastic Compute Cloud)

- AWS basic cloud web service
- A variety of instance types: general purpose, memory optimized, storage optimized, high performance, etc.
- Charges based on instance type, CPU, and memory allocation.
- t2.micro instances are in Free Tier

Environments on AWS (EC2 instances)

```
erebuss-MacBook-Pro-3:wikibase baskausj$ docker-machine ls
NAME          ACTIVE DRIVER   STATE URL                               SWARM   DOCKER   ERRORS
aws-sandbox   *      amazonec2 Running tcp://3.84.39.45:2376           v18.09.2
wikibase     -      amazonec2 Running tcp://[redacted]:2376           v18.09.2
erebuss-MacBook-Pro-3:wikibase baskausj$
```

The screenshot shows the AWS Management Console. At the top, there's a table of EC2 instances. The 'aws-sandbox' instance is highlighted. Below the table, the details for the 'aws-sandbox' instance are shown. Red arrows from the terminal output point to specific fields in the console:

- security group (change port settings here)**: Points to the 'Security groups' field, which is 'docker-machine'.
- machine name**: Points to the 'Instance ID' field, which is 'i-0228ed0ac14770517'.
- IP address**: Points to the 'IPv4 Public IP' field, which is '3.84.39.45'.

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP	IPv6	Key Name
aws-sandbox	i-0228ed0ac14770517	t2.micro	us-east-1a	running	2/2 check...	None	ec2-3-84-39-45.compute-1.amazonaws.com	3.84.39.45	-	aws-sandbox
biggerMetaphactor...	[redacted]	t2.large	us-east-1a	stopped		None		-	-	metaphactoryblazegraph
Dspace-DEV	[redacted]	t2.small	us-east-1b	stopped		None	[redacted].compute-1.amazonaws.com	[redacted]	-	dspace-server
Omeka-Dev	[redacted]	m3.large	us-east-1a	stopped		None	[redacted]	-	-	cba-digital
vr-c-iiif-env	[redacted]	t2.micro	us-east-1a	running	2/2 check...	None	[redacted].compute-1.amazonaws.com	[redacted]	-	
wikibase	[redacted]	t2.large	us-east-1a	running	2/2 check...	None	[redacted].compute-1.amazonaws.com	[redacted]	-	wikibase

Instance: i-0228ed0ac14770517 (aws-sandbox)		Public DNS: ec2-3-84-39-45.compute-1.amazonaws.com	
Description	Status Checks	Monitoring	Tags
Instance ID	i-0228ed0ac14770517		
Instance state	running		
Instance type	t2.micro		
Elastic IPs			
Availability zone	us-east-1a		
Security groups	docker-machine. view inbound rules. view outbound rules		
Scheduled events	No scheduled events		
AMI ID	ubuntu/images/hvm-ssd/ubuntu-xenial-16.04-amd64-server-20180228.1 (ami-927185ef)		
Platform	-		
IAM role	-		
Key pair name	aws-sandbox		
Owner	555751041262		
Launch time	February 26, 2019 at 9:12:29 AM UTC-6 (less than one hour)		
Termination protection	False		
Public DNS (IPv4)	ec2-3-84-39-45.compute-1.amazonaws.com		
IPv4 Public IP	3.84.39.45		
IPv6 IPs	-		
Private DNS	ip-172-31-86-55.ec2.internal		
Private IPs	172.31.86.55		
Secondary private IPs			
VPC ID	vpc-7234050a		
Subnet ID	subnet-4e940761		
Network interfaces	eth0		
Source/dest. check	True		
T2/T3 Unlimited	Disabled		
EBS-optimized	False		
Root device type	ebs		
Root device	/dev/sda1		

Demo

- Create an EC2 instance using Docker CLI:

```
docker-machine create --driver amazonec2 --amazonec2-open-port 8000 baskauf-sandbox
```

- Switch to the new environment (Linux only):

```
eval $(docker-machine env baskauf-sandbox)
```

- view environments (* means active for commands)

```
docker-machine ls
```

- pull and run an image

```
docker run -d -p 8000:80 --name webserver kitematic/hello-world-nginx
```

- dereference IP address in browser (port 8000)

Managing the server

- The server environment can be terminated from the EC2 web console or

`docker-machine rm name`

- In order to stop and restart the server, a static IP is required (otherwise the IP changes and daemon is lost)
- All kinds of problems (security certificates, mapping domain name, port mappings, etc.)
- You can SSH into Linux in the environment:

`docker-machine ssh baskauf-sandbox`

`sudo docker container ls -a`

`exit`

- See <https://docs.docker.com/machine/examples/aws/> for more info

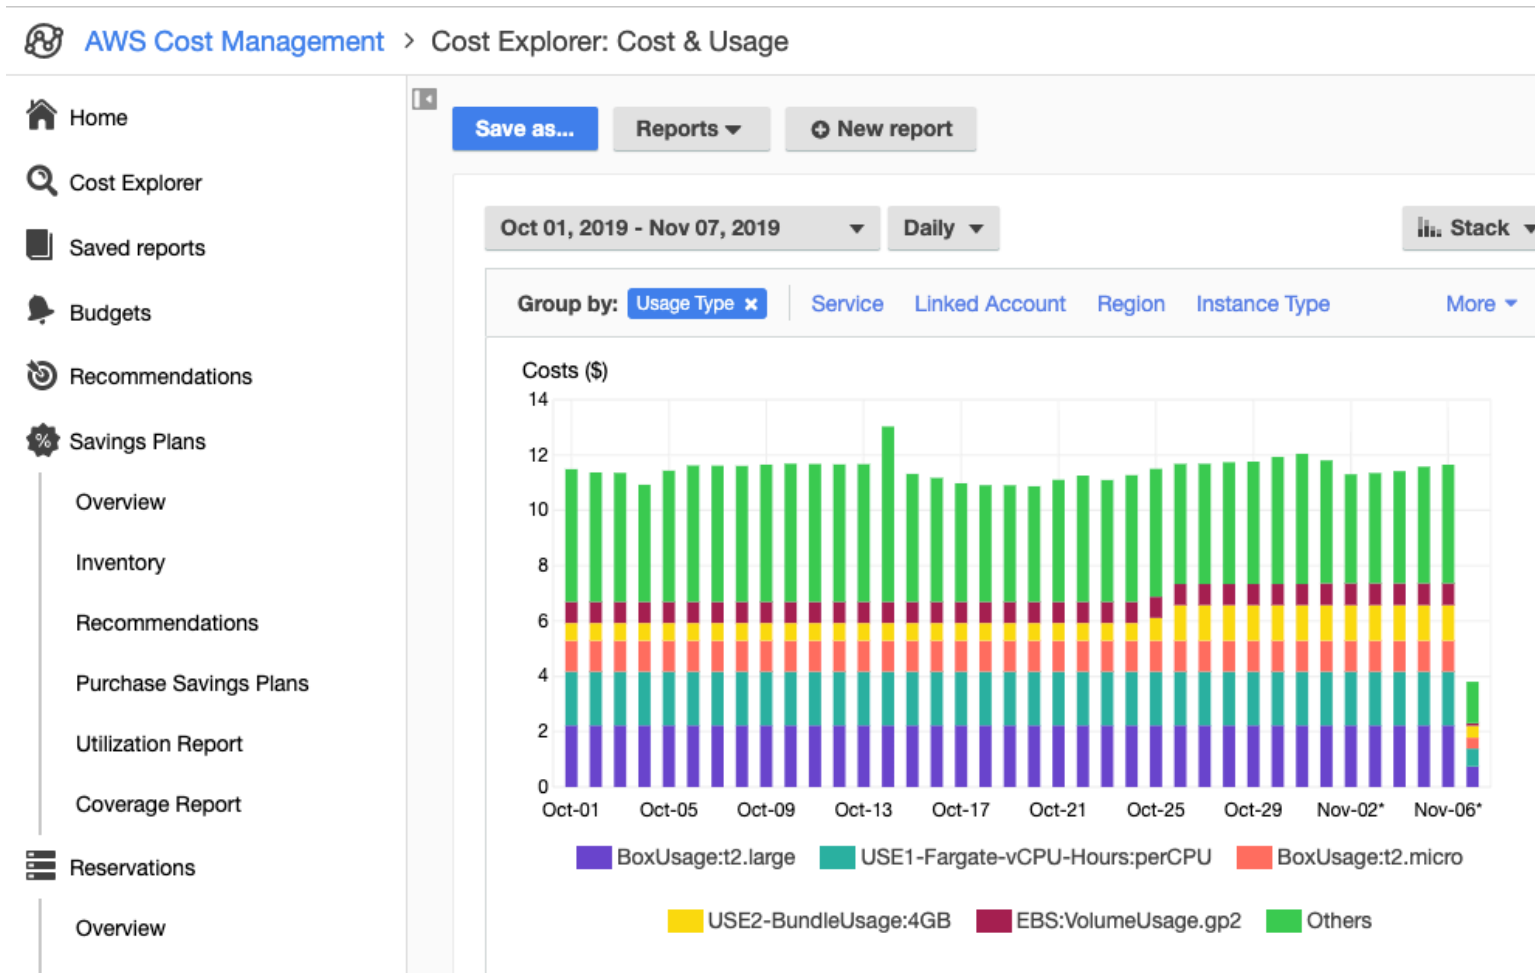
Variations built on EC2

- ECS (Elastic Container Service) = cluster of EC2 instances
- Elastic Beanstalk = operates on top of ECS. Automatically scales up by adding more EC2 instances as needed, plus load balancing, and security.
- Lightsail = simplified management, default security settings
- Services can pull container images directly from Docker Hub when you initiate them.

Cost Management tools


Cost Explorer

- t2.large instance of Wikibase is about \$2 per day.



Cost Management tools

- Budgets



Budget History (\$)

[View in AWS Cost Explorer](#) [Download CSV](#)

