



EMORY
LIBRARIES &
INFORMATION
TECHNOLOGY



Blair Boyd
Product Manager

Circe Tsui
Associate Director, Solutions Architecture

Agenda

- What is AWS?
- AWS at Emory
- Research Use Cases
- Demo
- Service Offering
- Questions

What is Amazon Web Services

- Market leader of Infrastructure as a Service (IaaS) providers
- Over 100 Products and Services
- No up-front investment
- Pay-as-you-go pricing
- Fast, elastic, scalable and with high availability



Amazon EC2
i.e. Virtual Computer



Amazon RDS
i.e. Relational Database Service



Amazon S3
i.e. “Box” on AWS

AWS at Emory

- In the near future, we will have an Emory-specific environment that our users should move to
- Until that time, we are assisting researchers set up an interim AWS account
- The interim account will require the use of a Pcard or credit card
- At this time, you cannot upload any PHI or sensitive data.

Pricing Examples as of March 2017

Resource	Capacity	Cost *
On-demand EC2 Server (t2.2xlarge) for General Purpose	8 Cores with 32 GiB	\$3 for 8 hours
On-demand EC2 Server (r4.8xlarge) for memory intensive applications	32 Cores with 244 GiB	\$17 for 8 hours
General purpose EBS Volume	100 GB	\$7/month
On-demand small Oracle database (db.t2.small)	100 GB	\$32.89 / month
S3 Storage standard	100 GB	\$2.18 / month
Glacier storage	100 GB	\$0.4 / month

* Data Transfer out and request fees may apply.

See <http://calculator.s3.amazonaws.com/> for details

A large, stylized, light gold torch graphic is positioned on the left side of the slide, extending from the top left towards the bottom center. The torch has a flame at the top and a handle with a small ring at the bottom. It is partially enclosed by a thick, curved gold line that starts near the top left and curves around the torch.

Research Use Cases

School of Nursing Project

- Needed a robust Oracle database to host 100,000 de-identified Emory Health Care patient records for both research and teaching. The data included hospital stays, patient visits, labs, diagnoses, etc.

AWS Relational Database Service

- Configure



Ready in 10 to 30 minutes

Locally Hosted

- Acquire a server
- Install Operating System
- Install Oracle Software
- Create a database
- Create a user




Ready in 5 days to 5 weeks

Acknowledgement: Andi Plotsky

Emory Integrated Genomic Core (EIGC) Project

- Needed to analyze 1000+ human genomes (~1 PB total) fast!
- With AWS, 200 genomes were analyzed in parallel with 200 EC2s, i.e. “virtual computers”, each with 32 vCPU and 244 GB memory.



Emory Integrated Genomic Core (EIGC) Project



Acknowledgement: Rich Johnston



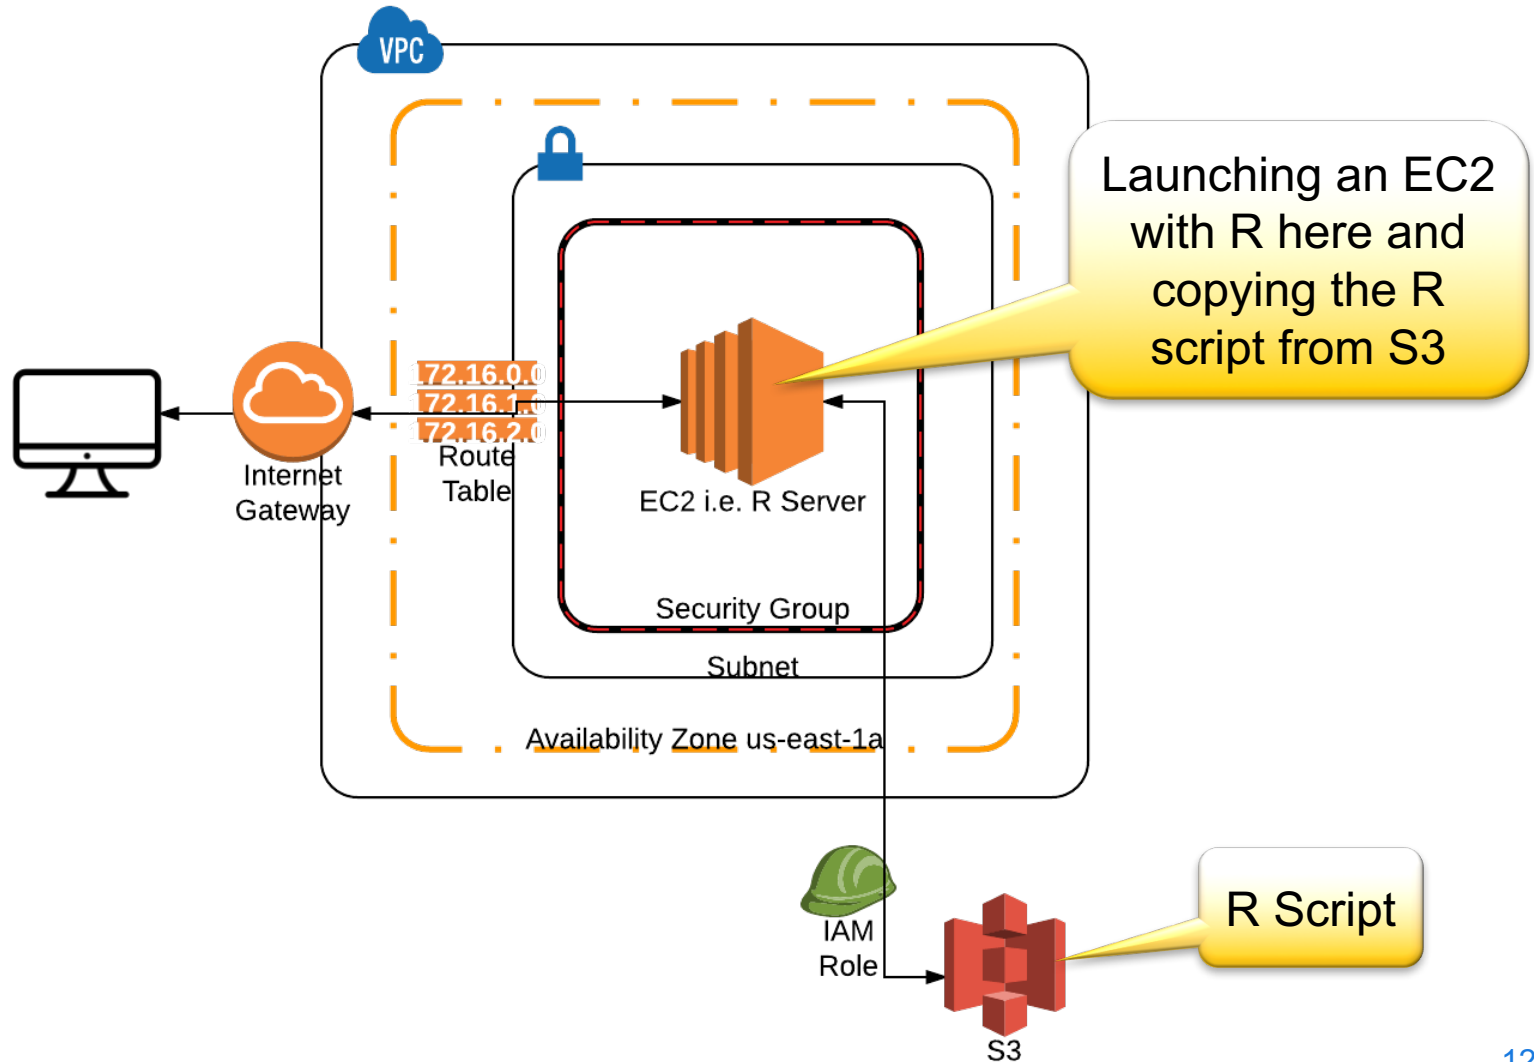
Demo

AWS Demo

Analyze 2.8 million records from the Centers for [Medicare & Medicaid Services](#) in an on-demand AWS m4.2xlarge (32G RAM / 8 vCPU) EC2 built with LITS's "R Server Amazon Machine Image" launched by a Cloudformation script.

Credits: <https://www.r-bloggers.com/bigglm-on-your-big-data-set-in-open-source-r-it-just-works-similar-as-in-sas/>

AWS Demo – R Server & the AWS Environment



A large, stylized, light yellow graphic of a torch with a flame, positioned diagonally on the left side of the slide. The torch is partially enclosed by a thick, curved yellow line that starts from the top left and curves downwards towards the bottom left.

Service Offering

How LITS Can Help:

- Setting up your AWS account
- Providing an Amazon Machine Image with commonly use scientific software packages
- Training
- Consultation by certified AWS Solution Architects
- Professional design, Implementation and Management under our service center model

Questions?

hpc.help@emory.edu

<https://wiki.service.emory.edu/display/public/Training>

Blair Boyd

- Email: blair.boyd@emory.edu
- Phone: 404.778.4916

Circe Tsui

- Email: circe.tsui@emory.edu
- Phone: 404.727.7035