



Banner in AWS: Technical Details

Gabriel Tocci

10-OCT-2018 @ 1:30pm - 2:30pm

Session ID # 1016



EAST TENNESSEE STATE
UNIVERSITY

Session Format

- 60 minute time slot
 - Touch on a lot of topics
 - Q&A anytime
- AWS / DBA
- Can you see REAL good?
- gabrieltozzi.com/talks



AWS Tech Summary

AWS services

- EC2 Instances
- Docker container orchestration
 - ECS
- Load balancing and Service discovery
 - ELB, ALB, HAProxy
- Continuous Delivery
 - Gitlab-ci, S3, ECR
- Networking
 - Route53, VPC
- Serverless functions
 - Lambda
- Databases
 - RDS: MS-Sql

Infrastructure as code (IaC)

- Terraform
- Puppet
- Docker

Monitoring

- Prometheus
- Alerts Manager
- Slack
- Graphana

Security

Disaster Recovery



Why Cloud?

- Efficiency and Scalability
 - Registration Spec'd Hardware
 - No need to guestimate capacity at time of purchase
- Self-Service Infrastructure
 - Dell, Cisco, Compellant, Who?
- Lower overall costs
 - Pay only for what you need
 - Cap-ex vs. op-ex
- Global resources
 - Global performance
 - DR
- Security First
- Elimination of hardware failures and upgrades
- Performance





There is no cloud
it's just someone else's computer



Why AWS?



- Market Leader
- Mature, robust, stable
- Support
- Well Documented
- Price
- CLI
- API
- Terraform Provider



Compute

Amazon EC2
Amazon Elastic Container Service
Amazon Elastic Container Service for Kubernetes
Amazon Elastic Container Registry
Amazon Lightsail
AWS Batch
AWS Elastic Beanstalk
AWS Fargate
AWS Lambda
AWS Serverless Application Repository
Auto Scaling
Elastic Load Balancing
VMware Cloud on AWS

Storage

Amazon Simple Storage Service (S3)
Amazon Elastic Block Storage (EBS)
Amazon Elastic File System (EFS)
Amazon Glacier
AWS Storage Gateway
AWS Snowball
AWS Snowball Edge
AWS Snowmobile

Database

Amazon Aurora
Amazon RDS
Amazon DynamoDB
Amazon ElastiCache
Amazon Redshift
Amazon Neptune
AWS Database Migration Service

Migration

AWS Migration Hub
AWS Application Discovery Service
AWS Database Migration Service

Networking & Content Delivery

Amazon VPC
Amazon CloudFront
Amazon Route 53
Amazon API Gateway
AWS Direct Connect
Elastic Load Balancing

Developer Tools

AWS CodeStar
AWS CodeCommit
AWS CodeBuild
AWS CodeDeploy
AWS CodePipeline
AWS Cloud9
AWS X-Ray
AWS Tools & SDKs

Management Tools

Amazon CloudWatch
AWS CloudFormation
AWS CloudTrail
AWS Config
AWS OpsWorks
AWS Service Catalog
AWS Systems Manager
AWS Trusted Advisor
AWS Personal Health Dashboard
AWS Command Line Interface
AWS Management Console
AWS Managed Services

Media Services

Amazon Elastic Transcoder
Amazon Kinesis Video Streams
AWS Elemental MediaConvert
AWS Elemental MediaLive
AWS Elemental MediaPackage
AWS Elemental MediaStore

Machine Learning

Amazon SageMaker
Amazon Comprehend
Amazon Lex
Amazon Polly
Amazon Rekognition
Amazon Machine Learning
Amazon Translate
Amazon Transcribe
AWS DeepLens
AWS Deep Learning AMIs
Apache MXNet on AWS
TensorFlow on AWS

Analytics

Amazon Athena
Amazon EMR
Amazon CloudSearch
Amazon Elasticsearch Service
Amazon Kinesis
Amazon Redshift
Amazon QuickSight
AWS Data Pipeline
AWS Glue

Security, Identity & Compliance

AWS Identity and Access Management (IAM)
Amazon Cloud Directory
Amazon Cognito
Amazon GuardDuty
Amazon Inspector
Amazon Macie
AWS Certificate Manager
AWS CloudHSM
AWS Directory Service
AWS Key Management Service
AWS Organizations
AWS Single Sign-On

AR & VR

Amazon Sumerian

Application Integration

Amazon MQ
Amazon Simple Queue Service (SQS)
Amazon Simple Notification Service (SNS)
AWS AppSync
AWS Step Functions

Customer Engagement

Amazon Connect
Amazon Pinpoint
Amazon Simple Email Service (SES)

Business Productivity

Alexa for Business
Amazon Chime
Amazon WorkDocs
Amazon WorkMail

Desktop & App Streaming

Amazon WorkSpaces
Amazon AppStream 2.0

Internet of Things

AWS IoT Core
Amazon FreeRTOS
AWS Greengrass
AWS IoT 1-Click
AWS IoT Analytics
AWS IoT Button
AWS IoT Device Defender
AWS IoT Device Management

Game Development

Amazon GameLift
Amazon Lumberyard

Software

ETSU AWS Project

- Current State is hybrid cloud / on prem
- Timing with Banner 9
 - Server Requisition
 - Technology Change
- 2 DBA, 1 manager/director
- 4 senior programmers, 2 junior programmer
- 0 system admin
- 0 network engineers



IaC

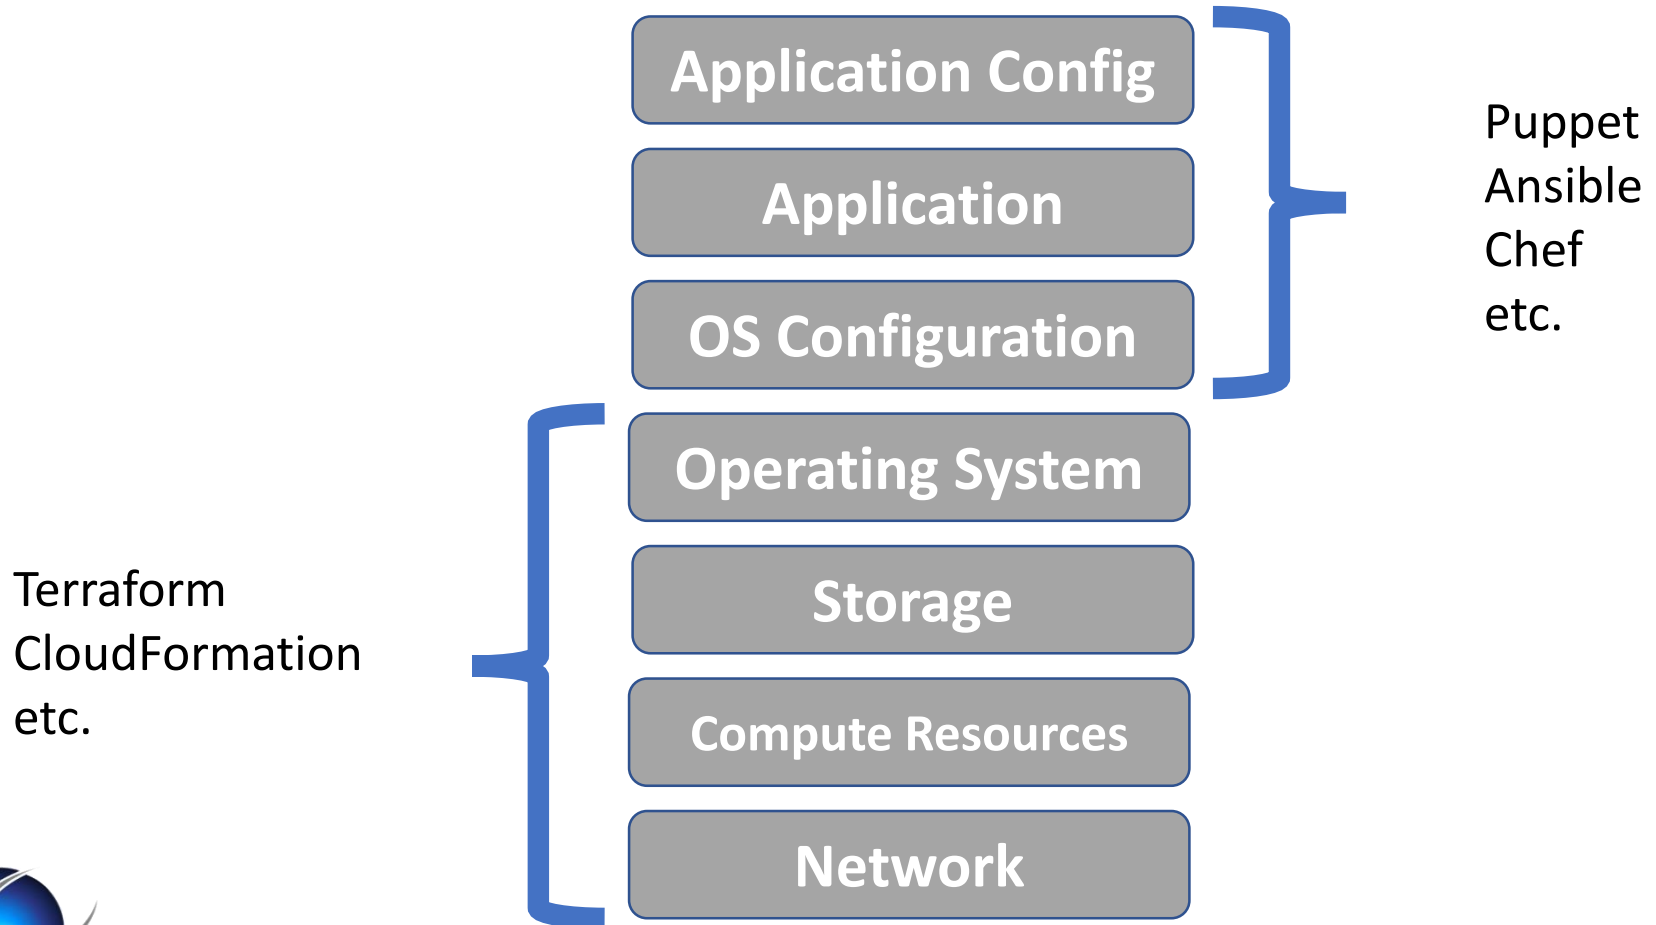
- Declares intent and interactions of resources
- Auditable infrastructure
- Increases recovery speed
- Reduces go-live errors
- Manages “drift” of configurations

Tools

- Puppet
- Terraform
- Docker
- Git



Provisioning vs. Configuration Management



Terraform

```
resource "aws_instance" "foo" {
  ami = "ami-ae7bfdb8"
  instance_type = "m4.xlarge"
  subnet_id = "${module.vpc.private_subnets[1]}"
  key_name = "linux_ec2"
  private_ip = "121.43.22.41"
  vpc_security_group_ids = ["${aws_security_group.dban.id}"]

  root_block_device{
    volume_type = "gp2"
    volume_size = "100"
    delete_on_termination = "false"
  }
}
```



Apps in AWS on EC2

- Oracle: Banner, ODS, Degreeworks, UC4
- Jobsub / UC4
- ESM
- Degreeworks: Classic Server and Rabbitmq

- Small Instances
 - Informatica Cloud Agent (salesforce / Targetx)
 - HAProxy
 - dtncommerce
 - dns



EC2: Instance Types





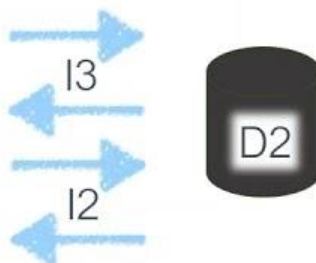
General purpose

Compute optimized

Storage and I/O optimized

Memory optimized

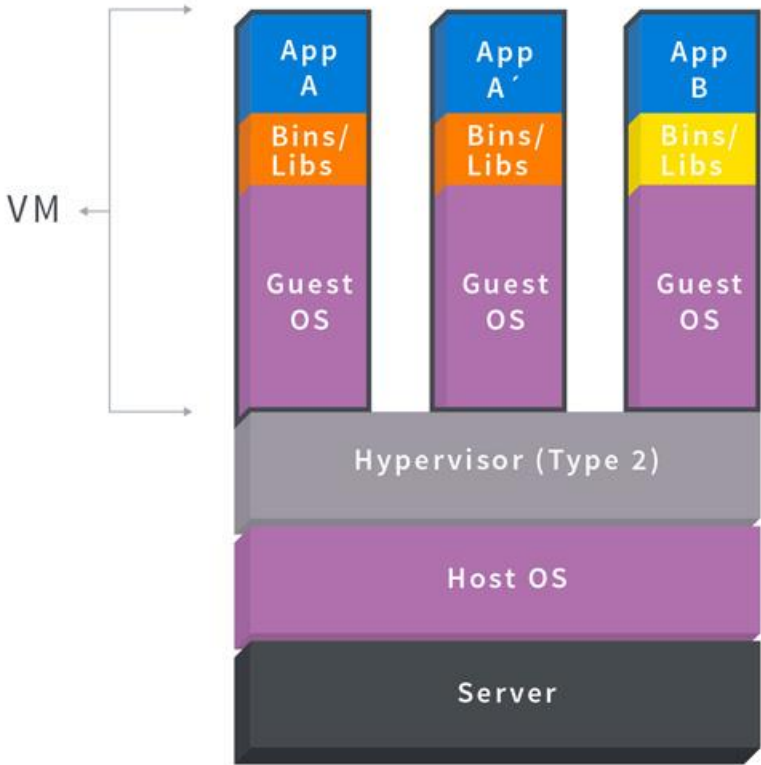
GPU or FPGA enabled



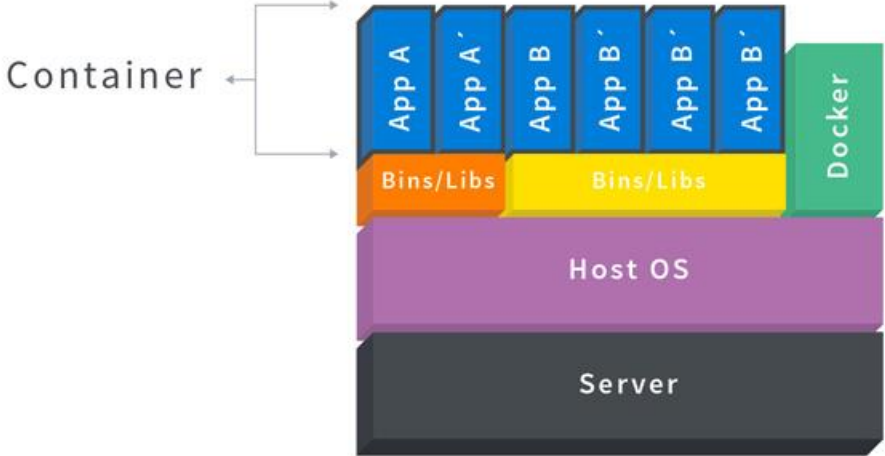
- <https://git.etsu.edu/aws/terraform/blob/master/development/dadb.tf>
- <https://calculator.s3.amazonaws.com/index.html>



Containers vs. VMs



Containers are isolated, but share OS and, where appropriate, bins/libraries



Orchestration: AWS Elastic Container Service (ECS)

- ECS Task Management
 - Desired Task Count
 - URL Routing
- EC2 Clusters
 - EC2 Hosts vs. Fargate
- Elastic Container Registry (ECR)
- Cloudwatch logs
 - 6 month retention



Containerization

- Git Repo's
 - Application
 - Docker: Base Images
 - ECS Service
- Externalized Configuration
 - Config Files
 - Docker.ini
 - Secrets.yaml
- Gitlab-ci.yaml
 - Push Tarball
 - Application and Configuration
- Entrypoint.py
 - Configure Application
 - Application Configuration
 - Launch Catalina



Admin Pages – Deployment Walkthrough

- Deploy via ESM to “dummy” host
- Update repo w/war or new configuration settings:
 - <https://git.etsu.edu/ecs/AdminPages/tree/master/tsth>
- GitlabCI executes `deploy_to_s3.py`
 - <https://git.etsu.edu/ecs/AdminPages/blob/master/.gitlab-ci.yml>
- `Deploy_to_s3.py` builds tarballs and pushes them to s3
 - https://git.etsu.edu/ecs/AdminPages/blob/master/deploy_to_s3.py
- Ecs task
 - <https://git.etsu.edu/aws/terraform/blob/master/env-development/ecs-tasks/tsth-BannerAdmin.json#L4>
- Tomcat image has `entrypoint.py`
 - <https://git.etsu.edu/docker/tomcat8-jre7-base/blob/master/entrypoint.py>
 - Configures war (docker.ini)
 - Configures secrets (secrets.yaml)
- Secrets.yaml get updated on password rotations
 - <https://git.etsu.edu/banner/password-fabric/blob/master/aws.py>

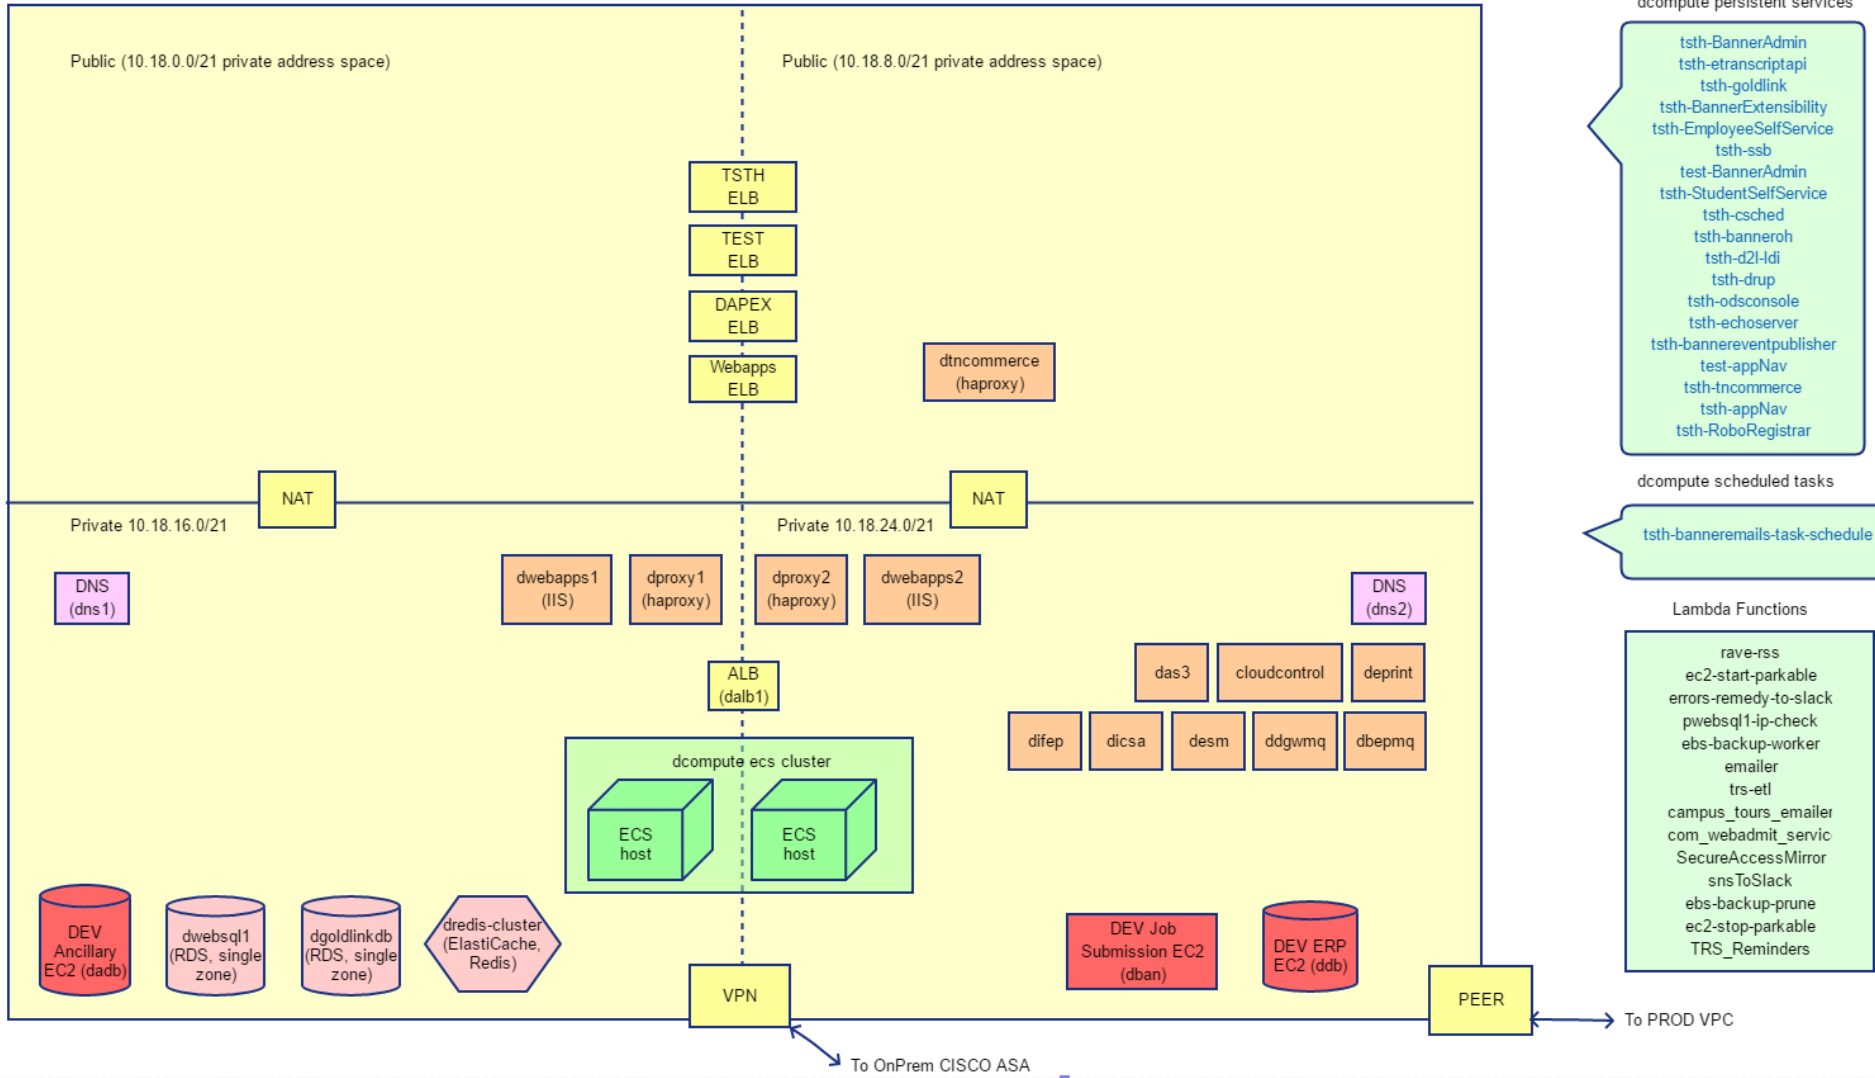


Containerized Banner Apps

- Application Navigator
- Admin Pages
- Employee Self Service
- eTranscript API
- Goldlink (drupal portal)
- Student Self Service
- College Scheduler
- Banner Online Help
- D2L – LDI
- General Self Service
- Integration API
- SSB
- Banner Event Publisher
- TNCommerce
- RoboRegistrar
- Banner Extensibility

<https://git.etsu.edu/aws/terraform/tree/master/env-development/ecs-tasks>





dcompute persistent services

- tsth-BannerAdmin
- tsth-etranscriptapi
- tsth-goldlink
- tsth-BannerExtensibility
- tsth-EmployeeSelfService
- tsth-ssb
- test-BannerAdmin
- tsth-StudentSelfService
- tsth-csched
- tsth-banneroh
- tsth-d2l-ldi
- tsth-drup
- tsth-odsconsole
- tsth-echoserver
- tsth-bannereventpublisher
- test-appNav
- tsth-tcommerce
- tsth-appNav
- tsth-RoboRegistrar

dcompute scheduled tasks

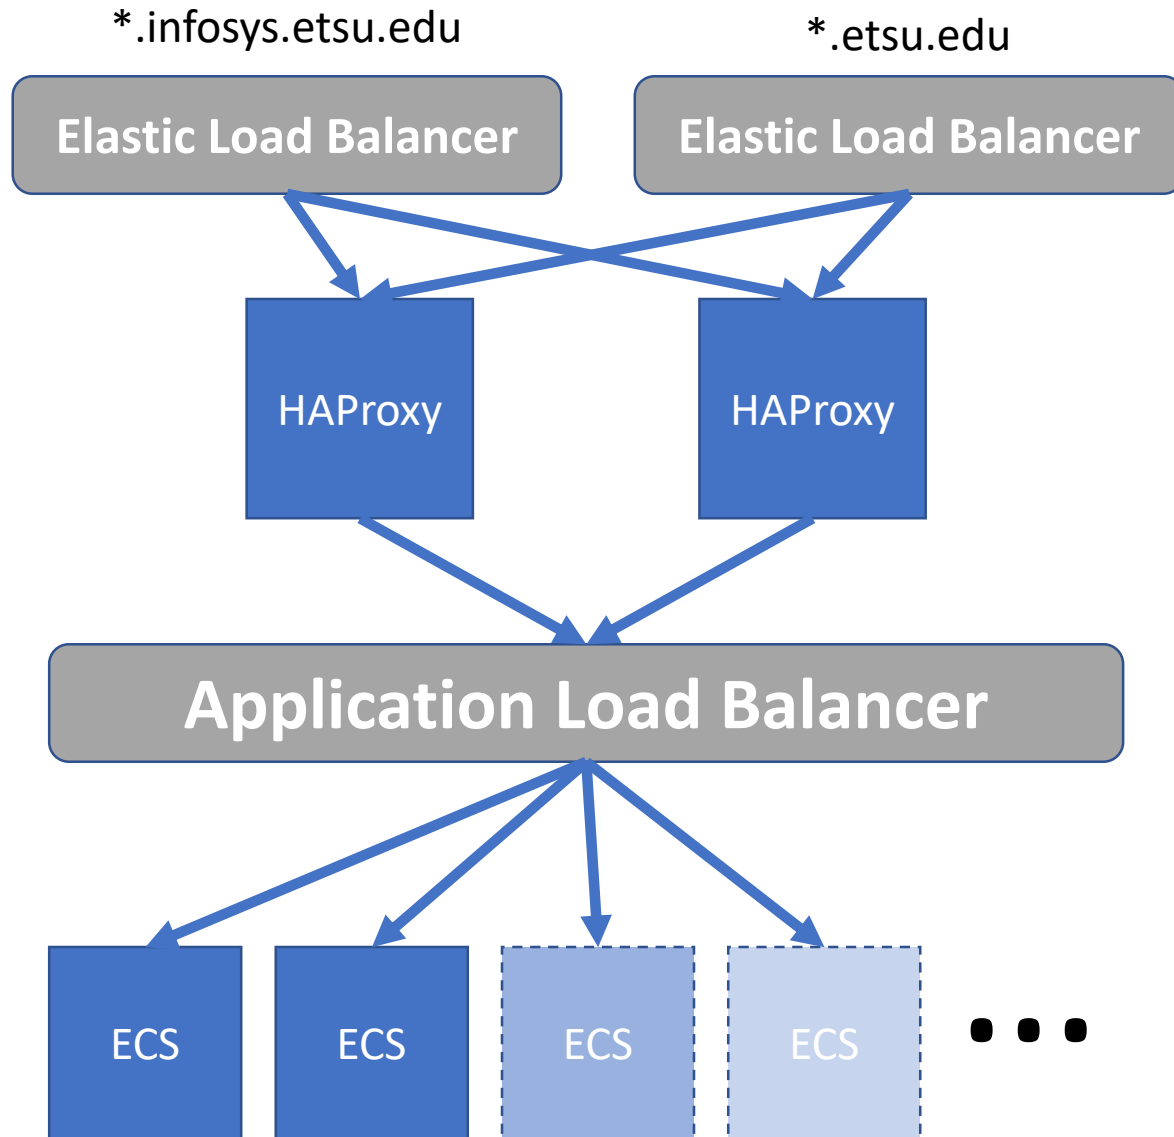
- tsth-banneremails-task-schedule

Lambda Functions

- rave-rss
- ec2-start-parkable
- errors-remedy-to-slack
- pwebsql1-ip-check
- ebs-backup-worker
- emailer
- trs-etl
- campus_tours_emailer
- com_webadmit_servic
- SecureAccessMirror
- snsToSlack
- ebs-backup-prune
- ec2-stop-parkable
- TRS_Reminders



ECS Service Discovery



On-Prem

- Evisions MAPS
- Intellicheck
- Portal (drupal)
- EIS / SSO Manager
- INB
- Dataguard
- BDM
- FSAAtlas
- Oracle Enterprise Manager



Lambda Functions

- rave-rss
- ec2-start-parkable
- errors-remedy-to-slack
- pwebsql1-ip-check
- ebs-backup-worker
- Emailer
- trs-etl
- campus_tours_emailer
- com_webadmit_service
- SecureAccessMirror
- snsToSlack
- ebs-backup-prune
- ec2-stop-parkable
- TRS_Reminders



Security

- IAM users
- 2FA
- Security Groups
 - Ingress / Egress Rules
- SSH Keypair login only
 - Same as On-Prem
 - Puppet
 - Iptables



Disaster Recovery

- Single Region
- Multi AZ
- Snapshots
- Declarative Infrastructure = Infrastructure is easily rebuilt
- Dataguard on premise



Banner Password Rotations / AWS

- Fabric
- AWS API
- <https://git.etsu.edu/banner/password-fabric/blob/master/aws.py>



Monitoring

- Cloudwatch, yuk!
- Prometheus
- Alerts Manager / Slack
- Graphana
- Enterprise Manager



CloudWatch
Dashboards
Alarms

ALARM

0

INSUFFICIENT

10

OK

50

Billing

Events

Rules

Event Buses

Logs

Metrics

Favorites

+ Add a dashboard

CloudWatch > Log Groups > tsth-BannerAdmin > us-east-1/tsth-BannerAdmin/5750c1ff-1115-4131-986b-b6129f66d450


Filter events

| Time (UTC -04:00) | Message |
|-------------------|---|
| 2018-09-13 | |
| 09:08:54 | 13-Sep-2018 13:08:54.751 INFO [main] org.apache.coyote.AbstractProtocol.init Initializing ProtocolHandler ["http-apr-8080"] |
| 09:08:54 | 13-Sep-2018 13:08:54.768 INFO [main] org.apache.coyote.AbstractProtocol.init Initializing ProtocolHandler ["ajp-apr-8009"] |
| 09:08:54 | 13-Sep-2018 13:08:54.773 INFO [main] org.apache.catalina.startup.Catalina.load Initialization processed in 1229 ms |
| 09:09:00 | 13-Sep-2018 13:09:00.083 INFO [main] org.apache.catalina.core.StandardService.startInternal Starting service Catalina |
| 09:09:00 | 13-Sep-2018 13:09:00.083 INFO [main] org.apache.catalina.core.StandardEngine.startInternal Starting Servlet Engine: Apache Tomcat/8.0.36 |
| 09:09:00 | 13-Sep-2018 13:09:00.120 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deploying web application archive |
| 09:09:21 | 13-Sep-2018 13:09:21.145 INFO [localhost-startStop-1] org.apache.jasper.servlet.TldScanner.scanJars At least one JAR was scanned for TLDs yet contained no TLDs. All files were scanned but no TLDs were found in them. Skipping unneeded JARs during scanning can improve startup time and JSP compilation time. |
| 09:09:21 | configuration: classpath:bannerHelp_configuration.groovy |
| 09:09:21 | configuration: file:/root/.grails/bannerHelp_configuration.groovy |
| 09:09:21 | log4j:WARN No appenders could be found for logger (org.codehaus.groovy.grails.commons.cfg.ConfigurationHelper). |
| 09:09:21 | log4j:WARN Please initialize the log4j system properly. |
| 09:09:21 | log4j:WARN See http://logging.apache.org/log4j/1.2/faq.html#noconfig for more info. |
| 09:09:24 | configuration: classpath:bannerHelp_configuration.groovy |
| 09:09:24 | configuration: file:/root/.grails/bannerHelp_configuration.groovy |
| 09:09:37 | 13-Sep-2018 13:09:37.308 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deployment of web application archive |
| 09:09:37 | 13-Sep-2018 13:09:37.315 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deploying web application archive |
| 09:09:37 | 13-Sep-2018 13:09:37.321 WARNING [localhost-startStop-1] org.apache.catalina.startup.SetContextPropertiesRule.begin [SetContextPropertiesRule] |
| 09:09:38 | 13-Sep-2018 13:09:38.036 INFO [localhost-startStop-1] org.apache.jasper.servlet.TldScanner.scanJars At least one JAR was scanned for TLDs yet contained no TLDs. All files were scanned but no TLDs were found in them. Skipping unneeded JARs during scanning can improve startup time and JSP compilation time. |
| 09:09:38 | 13-Sep-2018 13:09:38.045 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deployment of web application archive |
| 09:09:38 | 13-Sep-2018 13:09:38.046 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deploying web application archive |
| 09:10:08 | 13-Sep-2018 13:10:08.206 INFO [localhost-startStop-1] org.apache.jasper.servlet.TldScanner.scanJars At least one JAR was scanned for TLDs yet contained no TLDs. All files were scanned but no TLDs were found in them. Skipping unneeded JARs during scanning can improve startup time and JSP compilation time. |
| 09:10:27 | 13-Sep-2018 13:10:27.993 INFO [localhost-startStop-1] org.apache.catalina.startup.HostConfig.deployWAR Deployment of web application archive |
| 09:10:28 | 13-Sep-2018 13:10:28.001 INFO [main] org.apache.coyote.AbstractProtocol.start Starting ProtocolHandler ["http-apr-8080"] |
| 09:10:28 | 13-Sep-2018 13:10:28.017 INFO [main] org.apache.coyote.AbstractProtocol.start Starting ProtocolHandler ["ajp-apr-8009"] |
| 09:10:28 | 13-Sep-2018 13:10:28.028 INFO [main] org.apache.catalina.startup.Catalina.start Server startup in 93254 ms |




Basic General Info

Total responses by HTTP code



| | min | max | avg | current |
|----------------|-----|-----|-----|---------|
| Frontend 1xx | 0 | 0 | 0 | 0 |
| Frontend 2xx | 0 | 34 | 11 | 18 |
| Frontend 3xx | 0 | 3 | 0 | 0 |
| Frontend 4xx | 0 | 2 | 1 | 1 |
| Frontend 5xx | 0 | 0 | 0 | 0 |
| Frontend other | 0 | 0 | 0 | 0 |
| Backend 1xx | 0 | 0 | 0 | 0 |
| Backend 2xx | 0 | 15 | 2 | 1 |
| Backend 3xx | 0 | 0 | 0 | 0 |
| Backend 4xx | 0 | 0 | 0 | 0 |
| Backend 5xx | 0 | 0 | 0 | 0 |
| Backend other | 0 | 0 | 0 | 0 |

Current total of incoming / outgoing bytes



Total number of connections



Summary

- Take inventory of current resources – network, hardware, software, personnel
- Get Comfortable with AWS services
- Learn how to use the AWS cost estimator
- IaC
- cattle not pets



Whats Next?

- Containerize more apps
- Puppetize more vm configuration
- Create standby in separate region
- Increase usage of CI/CD
- Increase monitoring via Prometheus and alerting
- Greenfield
 - AWS Fargate
 - Kubernetes



Questions?



Resources

- <http://www.gabrieltocci.com/talks>
- BanDock: Docker Group
 - bandock@googlegroups.com
 - <https://bitbucket.org/edurepo/>
- BannerInTheCloud: AWS Group
 - <https://bannerinthecloud.slack.com>

