# Lessons Learned When a Dev Does Opsy Things

**Cora Fedesna**
(she/her)

@CorainChicago

noti.st/corainchicago

@DeliveryConf

# What is ActiveCampaign?

Automation. Email. CRM. Messaging. For **over 80,000** growing businesses.


ActiveCampaign >

# Software Engineer

# Developer

# Conference Organizer

@CorainChicago

noti.st/corainchicago

@DeliveryConf

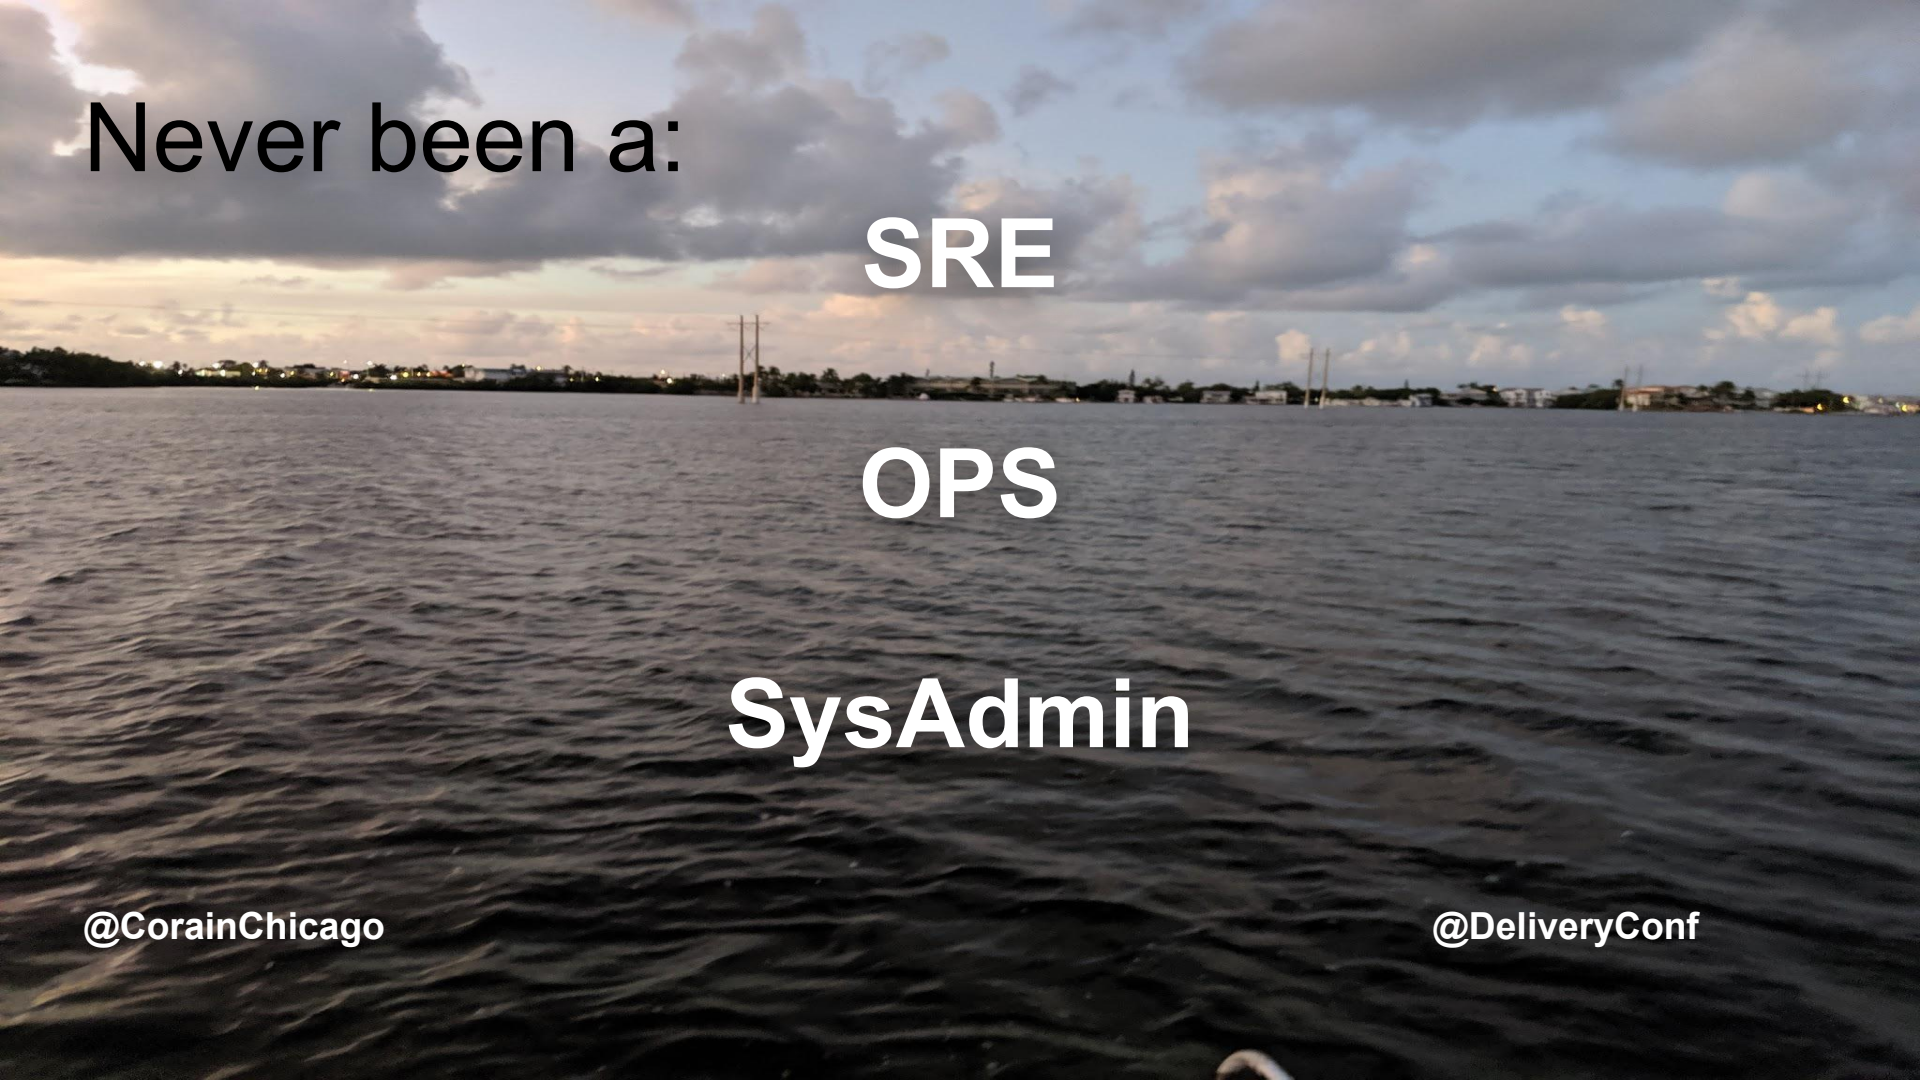
Never been a:
SRE
OPS
SysAdmin
@CorainChicago
@DeliveryConf

# Our story begins with dealignment of the SREs

(Also known as when I join the company)

# SRE Alignment

# SRE Dealignment


# We get to do opsy things


We have to do opsy things


@CorainChicago

@DeliveryConf


Jira


Segue / SEARCH-335


Jira


Segue / SEARCH-335


ASSIGNEE
Cora Fedesna

# Continuous Integration for your 1000% coverage

cfedesna — -bash — 80×24

```
Last login: Sun Dec 29 16:05:27 on console
Starting "default"...
Machine "default" is already running.
CFEDESNA-CHI:~ cfedesna$
```

# Continuous Integration Locally

```
cfedesna — -bash — 80×24
Last login: Sun Dec 29 16:05:27 on console
Starting "default"...
Machine "default" is already running.
CFEDESNA-CHI:~ cfedesna$
```

**gradle/mvn clean test**

# Continuous Deployment Locally

cfedesna — -bash — 80×24

```
Last login: Sun Dec 29 16:05:27 on console
Starting "default"...
Machine "default" is already running.
CFEDESNA-CHI:~ cfedesna$
```

**gradle/mvn clean build**

# Tasks

1. Write terraform to make my CI/CD pipeline

# Tasks

1. Write terraform to make my CI/CD pipeline
2. Write terraform to create the infrastructure for my application

# Tasks

1. Write terraform to make my CI/CD pipeline
2. Write terraform to create the infrastructure for my application
3. Get someone to apply the terraform

# Tasks

1. Write terraform to make my CI/CD pipeline
2. Write terraform to create the infrastructure for my application
3. Get someone to apply the terraform
4. Learn how to write terraform

# terraform/ci

- aws-main.tf
- build.yml
- codebuild-github-prs.tf
- variables.tf

## terraform/ci

- aws-main.tf
- build.yml
- codebuild-github-prs.tf
- variables.tf

## terraform/cd

- aws-main.tf
- build-docker-image.yml
- build-jar.yml
- deploy.yml
- pipeline-action-build-docker-image.tf
- pipeline-action-compile-jar.tf
- pipeline-action-deploy.tf
- pipeline.tf
- variables.tf

1. Tutorial
2. Code
3. PR
4. PR Comments - SRE
5. PR update
6. PR Comments - SRE
7. PR update
8. 1:1 with SRE
9. Code
10. PR
11. PR Comments - SRE
12. PR Update
13. PR Comments - SRE
14. PR Update
15. 1:1 Apply - SRE

# DEALIGNMENT

All teams are on a paging schedule.

Github

Gitlab

# Kubernetes

# .gitlab-ci.yml

```yaml
include:
  - project: 'devops/name-of-project'
    ref: v1.1.1
    file: 'filename.yml'
```

.gitlab-ci.yml

```
cache:
  paths:
    - .gradle/wrapper
    - .gradle/caches
```

.gitlab-ci.yml

```yaml
test:
  stage: test
  image: docker-image-for-jdk
  services:
    - docker:18-dind
  script:
    - export GRADLE_USER_HOME=`pwd`/.gradle
    - apk add docker
    - apk add bash
    - setup_docker
    - ./gradlew clean test
```

# .gitlab-ci.yml

```
build:
  stage: build
  image: image-jdk-used
  services:
    - docker:18-dind
  script:
    - export GRADLE_USER_HOME=`pwd`/.gradle
    - apk add docker
    - apk add bash
    - setup_docker
    - ecr_login name-of-aim-role
    - ./gradlew build
    - docker pull $REPOSITORY:latest || true
    - docker build --cache-from $REPOSITORY:latest --tag
$REPOSITORY:$TAG --tag $REPOSITORY:latest .
    - docker push $REPOSITORY:$TAG
    - docker push $REPOSITORY:latest
```

# .gitlab-ci.yml

```yaml
variables:
  VARIABLES_FROM_DEVOP_PROJECT_KEYS: value
  STAGING_ENABLED: "true"
  TEST_DISABLED: "true"
  REVIEW_DEPLOY_VALUES: >
    --env SPRING_PROFILES_ACTIVE=staging
    --env DEFAULT_LOG_LEVEL=debug
  STAGING_DEPLOY_VALUES: >
    --env SPRING_PROFILES_ACTIVE=staging
    --env DEFAULT_LOG_LEVEL=debug
```

# app.yaml

```yaml
name: service-name

namespace: service-name

image: service-name:latest

imagePullPolicy: IfNotPresent

http: true

gateway: service-name.test

iam: k8s/segue-e2e-role

containerPort: 8080

healthPath: /actuator/health

minScale: 1

maxScale: 1

livenessDelay: 360

livenessTimeout: 15

readinessDelay: 360

readinessTimeout: 50

memoryRequest: 2Gi

memoryLimit: 4Gi

vaultAddr:
https://vault.staging.app-us1.com:8200

secrets: ["vault/path/to/secret"]

prometheusPath: /actuator/prometheus
```

app.yaml

```
name: service-name

namespace: service-name

image: service-name:latest

containerPort: 8080

healthPath: /actuato

minScale: 1

maxScale: 1
```

```
livenessDelay: 360

livenessTimeout: 15

readinessDelay: 360

https://vault.staging.app-us1.com:8200

secrets: ["vault/path/to/secret"]

prometheusPath: /actuator/prometheus
```

Oops! Our devops team abstracted this work away from us.

# Book Club Time!

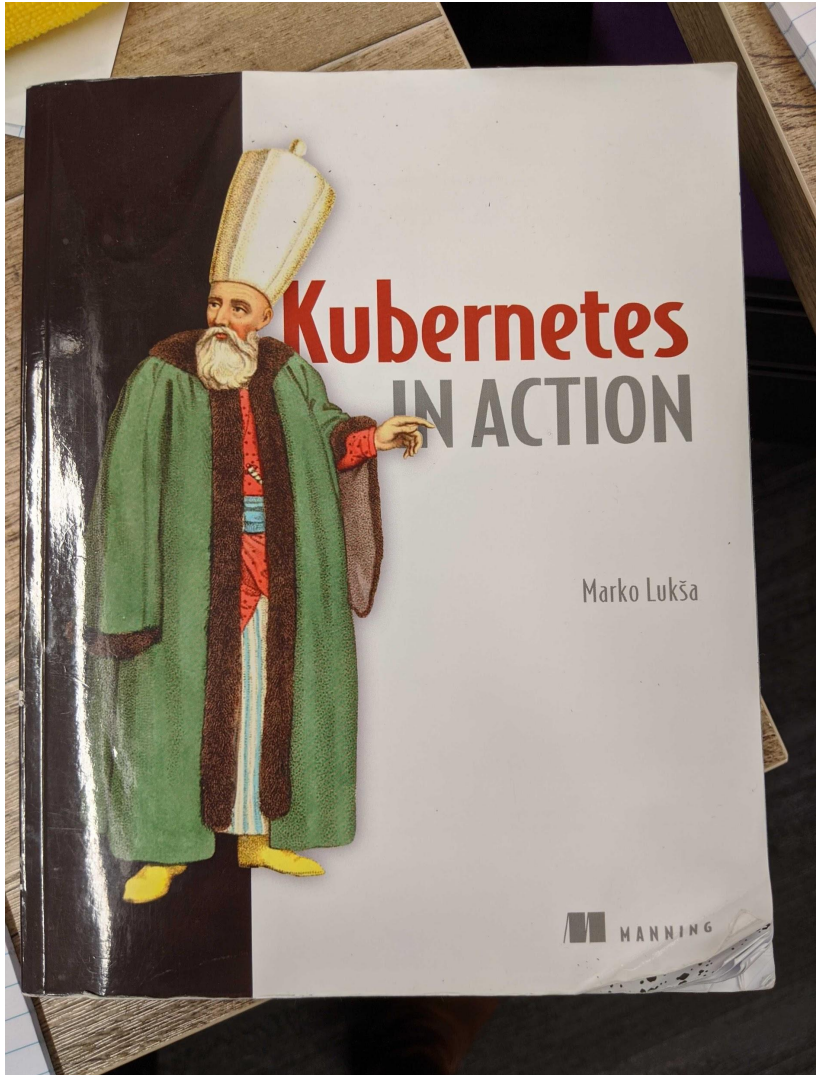

# Book Club Time!

(actually reading chapters 1-7)

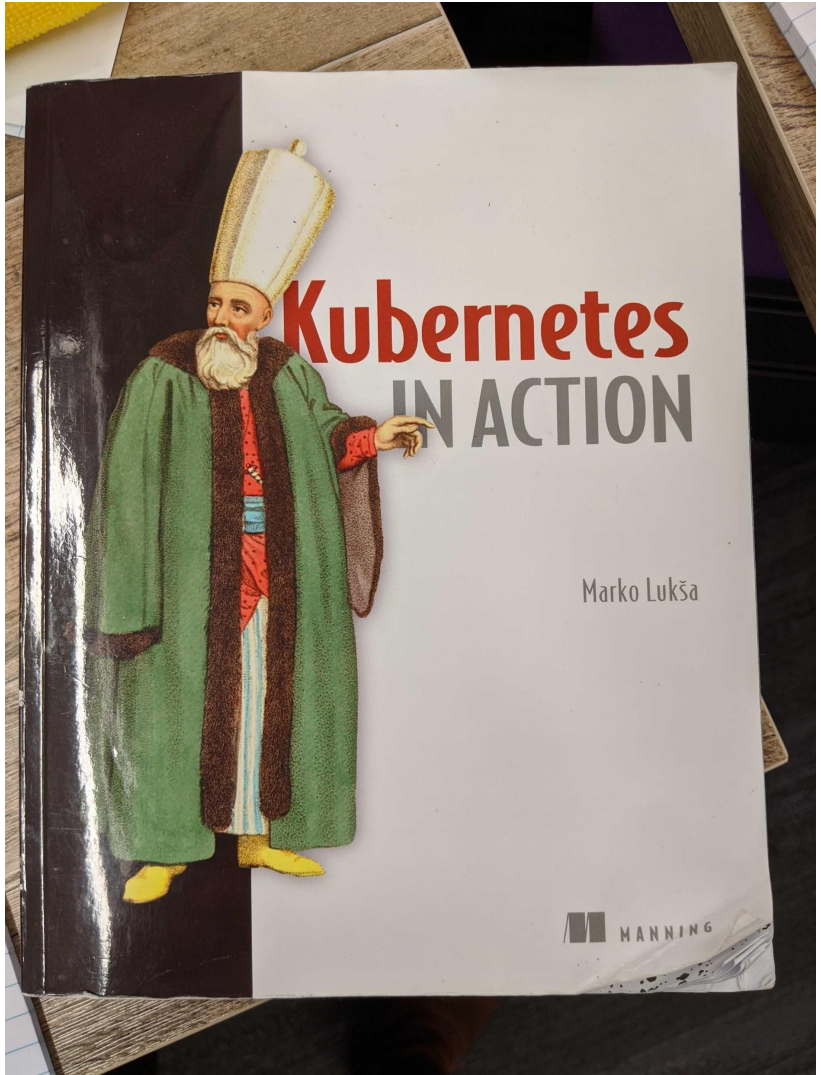

# Kubernetes


CI/CD being much easier, let’s go back to Terraform

# Terraform

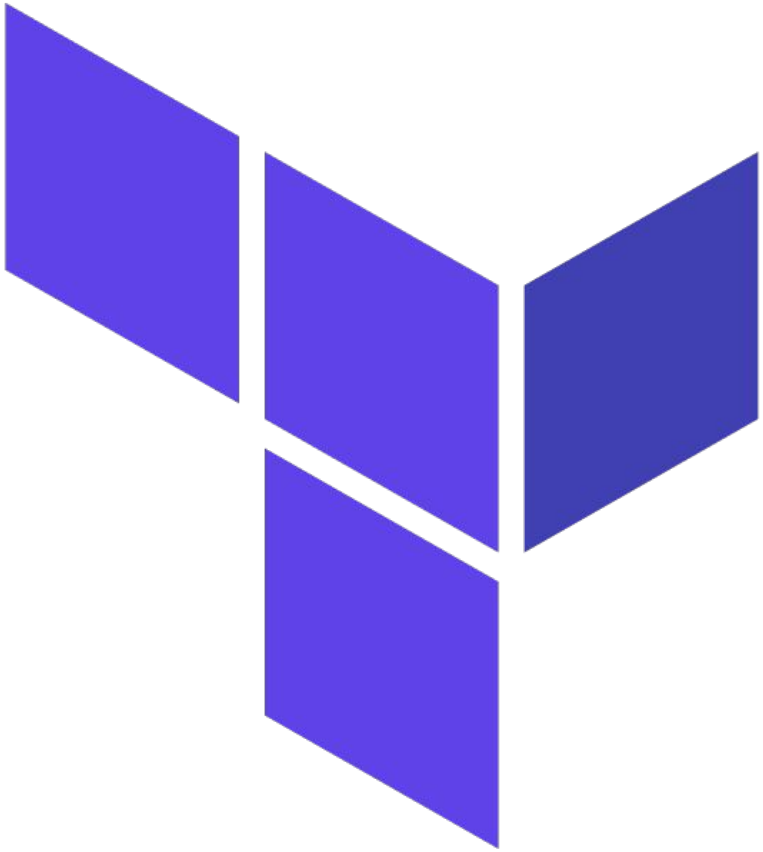

# Learn about provisioning infrastructure with HashiCorp Terraform

Get Started

Skip to Operations and Development Tracks

## Getting Started - AWS

12 TOPICS 69 MINS

In this introduction to Terraform, you will learn how to use Terraform to build infrastructure in AWS.

**Introduction to Infrastructure as Code with Terraform** →

5 MIN | What is Infrastructure as Code and Why is Terraform Useful?

**Installing Terraform** →

# Playing with AWS == Scary

**Warning!** If you're not using an account that qualifies under the AWS free-tier, you may be charged to run these examples. The most you should be charged should only be a few dollars, but we're not responsible for any charges that may incur.

# Dev View

Version

Username

Password

```
resource "aws_db_instance" "name_mysql" {
  instance_class       = "${var.mysql_instance_class}"
  db_subnet_group_name = "${aws_db_subnet_group.mysql.name}"
  engine               = "mysql"
  allocated_storage    = "${var.allocated_storage}"
  engine_version       = "${var.mysql_version}"
  username             = "${var.mysql_user_name}"
  password             = "${var.mysql_password}"

  vpc_security_group_ids = [
    "${data.aws_security_group.rds.id}",
  ]

  skip_final_snapshot = true
  identifier          = "${var.mysql_idenitifier}"

  tags = "${merge(
        local.common_tags,
        map(
          "Name", "${var.application_name}-db-instance"
        )
    )}"

  lifecycle {
    ignore_changes = [
      "password",
    ]
  }
}
```

# CODE

Good

Bad

# Terraform Coding == Java Coding

## Good

- Modular
- Variables used
- Documentation
- Examples available

## Bad

- All in one
- Hard code values
- Don’t follow examples

# Terraform Coding == Java Coding

## Good

- Modular
- Variables used
- Documentation
- Examples

## Bad

- All in one
- Hard code values
- Don’t follow examples
- Doesn’t run

# Terraform Coding == Java Coding

## Good

- Modular structure
- Variables used
- Documentation
- Examples

## Bad

- All in one
- Hard code values
- Don’t follow examples
- Doesn’t run
- Runs and does something inadvertent and horrible

# Module

```
module "service name" {
  source                = "../modules/filename"
  environment_type      = "staging"
  mysql_instance_class = "${var.mysql_instance_class}"
  mysql_version         = "${var.mysql_version}"
  mysql_password        = "${var.mysql_password}"
  mysql_user_name       = "${var.mysql_user_name}"
}
```

# Resources

```
resource "aws_s3_bucket" "bucket name" {
  bucket        = "${var.application_name}-controls"
  acl           = "private"
  force_destroy = true

  tags = "${merge(
    local.common_tags,
    map(
      "Name", "${var.application_name}-controls"
    )
  )}"
}
```

# Resources

```
resource "a" "" {}
```

```
resourc
```

- aws_vpc
- aws_subnet
- aws_db_instance
- aws_sns_topic
- aws_security_group
- aws_ecr_repository
- aws_eip
- aws_instance
- aws_codebuild_project
- acme_certificate
- acme_registration
- alicloud_actiontrail

Press ^. to choose the selected (or first) suggestion and insert a dot afterwards >>

# Resource: aws_s3_bucket

JUMP TO SECTION

Provides a S3 bucket resource.

## Example Usage

### Private Bucket w/ Tags

```
resource "aws_s3_bucket" "b" {
  bucket = "my-tf-test-bucket"
  acl    = "private"

  tags = {
    Name        = "My bucket"
    Environment = "Dev"
  }
}
```

# Data Blocks

```
data "aws_vpc" "my_name_for_it" {
  tags {
    Name = "region I want to filter by"
  }
}
```

# Data Blocks

```
data "aws_security_group" "rds"
{
  name = "rds"

  tags {
    Zone = "zone name"
  }
}
```

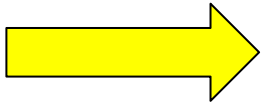

```
resource "aws_db_instance" "name_mysql" {
  instance_class       = "${var.mysql_instance_class}"
  db_subnet_group_name =
"${aws_db_subnet_group.mysql.name}"
  engine               = "mysql"
  allocated_storage    = "${var.allocated_storage}"
  engine_version       = "${var.mysql_version}"
  username             = "${var.mysql_user_name}"
  password             = "${var.mysql_password}"

  vpc_security_group_ids = [
    "${data.aws_security_group.rds.id}",
  ]
}
```

# Variables

```
variable "environment_type" {
  default     = "staging"
  description = "notes about it, explanation"
}

variable "mysql_password" {
  description = "The password for the mysql database"
  type        = "string"
}
```

# Data vs. Variable

# terraform.tfvars

```
mysql_instance_class = "db.m5.12xlarge"

mysql_user_name = "username"

kubernetes_vpc_name = "vpc_name"
```

terraform mysql

All | Videos | News | Images | Shopping | More | Settings | Tools

About 1,390,000 results (0.33 seconds)

## Provider: MySQL - Terraform by HashiCorp

https://www.terraform.io › docs › providers › mysql

**MySQL** is a relational database server. The **MySQL** provider exposes resources used to manage the configuration of resources in a **MySQL** server. Use the ...

### Mysql_user

Creates and manages a user on a MySQL server.

### Azurerm_mysql_database

Manages a MySQL Database within a MySQL Server ...

### Mysql_database

Caution: The mysql_database resource can completely delete ...

### Mysql_user_password

Creates and manages the password for a user on a ...

### Mysql_grant

The mysql_grant resource creates and manages privileges given ...

### Azurerm_mysql_server

resource "azurerm_resource_group ...

More results from terraform.io »

**Caution:** The `mysql_database` resource can completely delete your database just as easily as it can create it. To avoid costly accidents, consider setting `prevent_destroy` on your database resources as an extra safety measure.

## Example Usage

```
resource "mysql_database" "app" {
  name = "my_awesome_app"
}
```

## Argument Reference

The following arguments are supported:

- `name` - (Required) The name of the database. This must be unique within a given MySQL server and may or may not be case-sensitive depending on the operating system on which the MySQL server is running.
- `default_character_set` - (Optional) The default character set to use when a table is created without specifying an explicit character set. Defaults to "utf8".
- `default_collation` - (Optional) The default collation to use when a table is created without specifying an explicit collation. Defaults to `utf8_general_ci`. Each character set has its own set of collations, so changing the character set requires also changing the collation.

# Terraform - Tricky Things

**Destroy database**

**Commits passwords**

# Terraform - Tricky Things

Destroy database

Commits passwords

DEAR LORD - IT CAN FEEL SIDEWAYS


```
terraform fmt
```

Commit code

```
terraform init

terraform plan

terraform apply
```

@CorainChicago

@DeliveryConf

1. ~~Tutorial~~
2. Code
3. PR
4. PR Comments - SRE
5. PR update
6. ~~PR Comments - SRE~~
7. ~~PR update~~
8. 1:1 with SRE (briefer)
9. Code
10. PR
11. PR Comments - SRE
12. ~~PR Update~~
13. ~~PR Comments - SRE~~
14. PR Update
15. 1:1 Apply - SRE


DEALIGNMENT
My team's SRE stops coming to stand up.

Java

Spring

Docker

Java

Spring

Docker

Database

S3

API security

K8 Cluster

# What’s still a struggle?

# Things Devs Rarely Need to Think About

IAM Policies

Cidr Blocks

VPCs

Things Devs Rarely Deal With

# SECURITY

| | |
|---|---|
| **IAM** | Identity and Access Management |
| **Cidr Blocks** | Classless Inter-Domain Routing |
| **VPC** | Virtual Private Cloud |

# **IAM -** Identity and Access Management

An entity that, when attached to an identity or resource, defines their permissions.

AWS Permissions Thing

# IAM

Identity and Access Management

```
data "aws_iam_policy_document" "policy_name" {
  statement {
    effect = "Allow"
    actions = [
      "s3:PutObject",
      "s3:GetObject"
    ]
    resources = [
      "${aws_s3_bucket.name.arn}/*",
    ]
  }
}
```

# IAM - Identity and Access Management

```
data "aws_iam_policy_document" "policy_name" {

  statement {

    effect = "Allow"

    actions = [

      "s3:PutObject",

      "s3:GetObject",

    ]

    resources = [

      "${aws_s3_bucket.name.arn}/*",

    ]

  }

}
```

```
resource "aws_iam_policy" "name_task_policy" {

  name   = "name-task-profile"

  policy = "${data.aws_iam_policy_document.policy_name.json}"

}


resource "aws_iam_policy_attachment" "attachment_policy" {

  name       = "name-task-profile-attachment"

  roles      = ["${aws_iam_role.name_role.name}"]

  policy_arn = "${aws_iam_policy.name_task_policy.arn}"

}
```

# **CIDR -** Classless Inter-Domain Routing

192.168.100.14/26

# CIDR - Classless Inter-Domain Routing

| CIDR | Subnet mask (decimal) | Subnet mask (binary) | Available addresses | |
|---|---|---|---|---|
| /0 | 0.0.0.0 | 00000000.00000000.00000000.00000000 | 4.294.967.296 | $2^{32}$ |
| /1 | 128.0.0.0 | 10000000.00000000.00000000.00000000 | 2.147.483.648 | $2^{31}$ |
| /2 | 192.0.0.0 | 11000000.00000000.00000000.00000000 | 1.073.741.824 | $2^{30}$ |
| /3 | 224.0.0.0 | 11100000.00000000.00000000.00000000 | 536.870.912 | $2^{29}$ |
| /4 | 240.0.0.0 | 11110000.00000000.00000000.00000000 | 268.435.456 | $2^{28}$ |
| /5 | 248.0.0.0 | 11111000.00000000.00000000.00000000 | 134.217.728 | $2^{27}$ |
| /6 | 252.0.0.0 | 11111100.00000000.00000000.00000000 | 67.108.864 | $2^{26}$ |
| /7 | 254.0.0.0 | 11111110.00000000.00000000.00000000 | 33.554.432 | $2^{25}$ |
| /8 | 255.0.0.0 | 11111111.00000000.00000000.00000000 | 16.777.216 | $2^{24}$ |
| /9 | 255.128.0.0 | 11111111.10000000.00000000.00000000 | 8.388.608 | $2^{23}$ |
| /10 | 255.192.0.0 | 11111111.11000000.00000000.00000000 | 4.194.304 | $2^{22}$ |
| /11 | 255.224.0.0 | 11111111.11100000.00000000.00000000 | 2.097.152 | $2^{21}$ |
| /12 | 255.240.0.0 | 11111111.11110000.00000000.00000000 | 1.048.576 | $2^{20}$ |
| /13 | 255.248.0.0 | 11111111.11111000.00000000.00000000 | 524.288 | $2^{19}$ |
| /14 | 255.252.0.0 | 11111111.11111100.00000000.00000000 | 262.144 | $2^{18}$ |
| /15 | 255.254.0.0 | 11111111.11111110.00000000.00000000 | 131.072 | $2^{17}$ |
| /16 | 255.255.0.0 | 11111111.11111111.00000000.00000000 | 65.536 | $2^{16}$ |
| /17 | 255.255.128.0 | 11111111.11111111.10000000.00000000 | 32.768 | $2^{15}$ |
| /18 | 255.255.192.0 | 11111111.11111111.11000000.00000000 | 16.384 | $2^{14}$ |
| /19 | 255.255.224.0 | 11111111.11111111.11100000.00000000 | 8.192 | $2^{13}$ |

# VPC - Virtual Private Cloud

logically isolated
section of ~~the AWS~~
any cloud

**(Has a cidr block)**

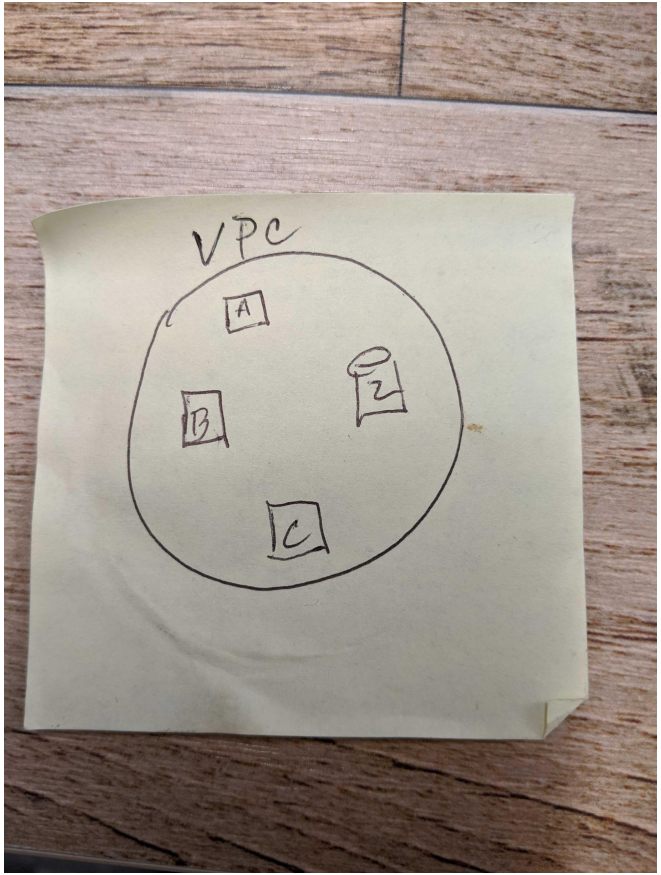

```
data "aws_vpc" "kubernetes" {
  tags {
    Name = "${var.kubernetes_vpc_name}"
  }
}
```

```
cidr_blocks = ["${data.aws_vpc.kubernetes.cidr_block}"]
```

| | |
|---|---|
| **IAM** | Identity and Access Management |
| **Cidr Blocks** | Classless Inter-Domain Routing |
| **VPC** | Virtual Private Cloud |

# TUNING

@CorainChicago

@DeliveryConf

# Generate Load

Volume hitting service?

Services falling over?

Database Connections - good?

Error logs?

Volum

er?

Con

# Will my service cause an alert or page?

Volun

# Can I prevent it?

Database
Connections - good?

Error logs?

# AWS Console

# Datadog Dashboards

# Kibana Logs

# Kubectl Commands

History

Simple Queue Service

Simple Notification Service

Billing

Console Home

Amazon DocumentDB

EC2

Find a service by name or feature (for example, EC2, S3 or VM, storage).

Group | A-Z

**Compute**
- EC2
- Lightsail
- ECR
- ECS
- EKS
- Lambda
- Batch
- Elastic Beanstalk
- Serverless Application Repository
- EC2 Image Builder

**Storage**
- S3
- EFS
- FSx
- S3 Glacier
- Storage Gateway
- AWS Backup

**Database**
- RDS
- DynamoDB

**Customer Enablement**
- AWS IQ
- Support
- Managed Services

**Blockchain**
- Amazon Managed Blockchain

**Satellite**
- Ground Station

**Quantum Technologies**
- Amazon Braket

**Management & Governance**
- AWS Organizations
- CloudWatch
- AWS Auto Scaling
- CloudFormation
- CloudTrail
- Config

**Analytics**
- Athena
- EMR
- CloudSearch
- Elasticsearch Service
- Kinesis
- QuickSight
- Data Pipeline
- AWS Data Exchange
- AWS Glue
- AWS Lake Formation
- MSK

**Security, Identity, & Compliance**
- IAM
- Resource Access Manager
- Cognito
- Secrets Manager
- GuardDuty
- Inspector
- Amazon Macie
- AWS Single Sign-On
- Certificate Manager
- Key Management Service

**Business Applications**
- Alexa for Business
- Amazon Chime
- WorkMail

**End User Computing**
- WorkSpaces
- AppStream 2.0
- WorkDocs
- WorkLink

**Internet Of Things**
- IoT Core
- Amazon FreeRTOS
- IoT 1-Click
- IoT Analytics
- IoT Device Defender
- IoT Device Management
- IoT Events
- IoT Greengrass
- IoT SiteWise
- IoT Things Graph

close

# Where’d my db go?

aws Services Resource Groups Engineer/cfedesna@activecam... N. Virginia Support

Database

## Amazon RDS
Managed relational database service

Easily set up, operate, and scale a relational database in the cloud.

**Get started**

Create a new database instance by selecting a database engine with your desired configuration.

Create database

Restore a database from Amazon S3

**Pricing and costs**

- Amazon Aurora Pricing
- Amazon RDS for MySQL Pricing
- Amazon RDS for PostgreSQL Pricing
- Amazon RDS for MariaDB Pricing
- Amazon RDS for Oracle Pricing
- Amazon RDS for SQL Server Pricing

**Getting started and documentation**

How it works

# Where’d my db go?

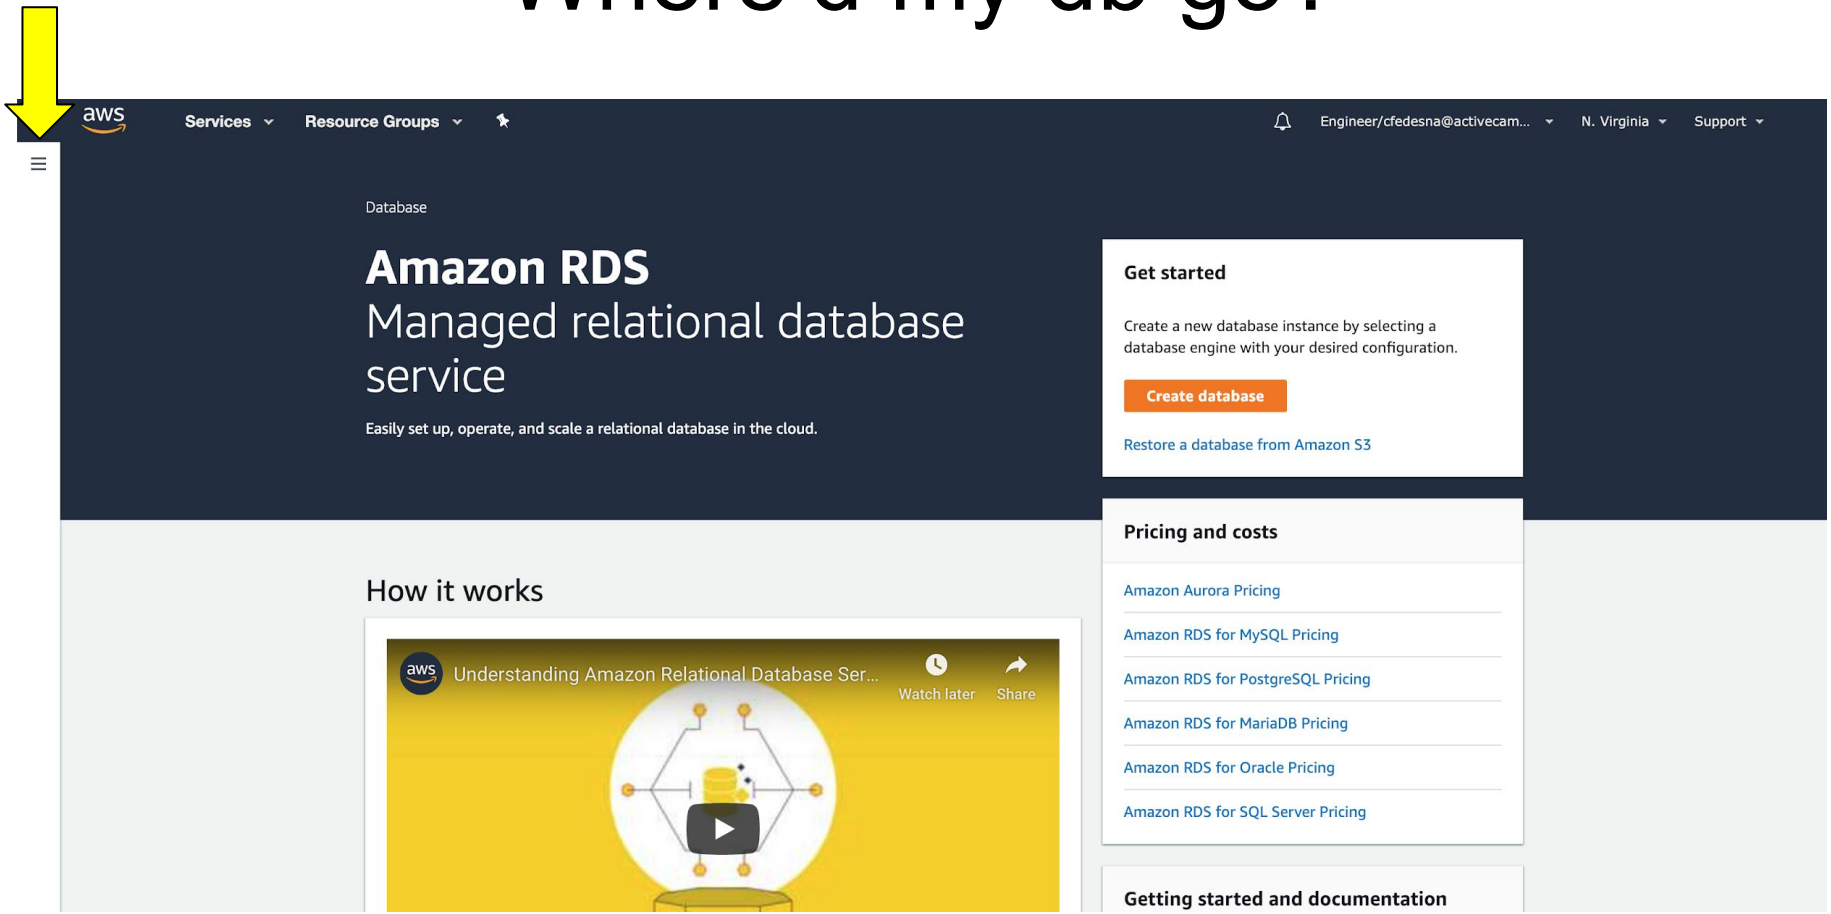

Find a service by name or feature (for example, EC2, S3 or VM, storage).

Group | A-Z

## Compute
- EC2
- Lightsail
- ECR
- ECS
- EKS
- Lambda
- Batch
- Elastic Beanstalk
- Serverless Application Repository
- AWS Outposts
- EC2 Image Builder

## Storage
- S3
- EFS
- FSx
- S3 Glacier
- Storage Gateway
- AWS Backup

## Database
- RDS

## Customer Enablement
- AWS IQ
- Support
- Managed Services

## Blockchain
- Amazon Managed Blockchain

## Satellite
- Ground Station

## Quantum Technologies
- Amazon Braket

## Management & Governance
- AWS Organizations
- CloudWatch
- AWS Auto Scaling
- CloudFormation
- CloudTrail
- Config
- OpsWorks

## Machine Learning
- Amazon SageMaker
- Amazon CodeGuru
- Amazon Comprehend
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Kendra
- Amazon Lex
- Amazon Machine Learning
- Amazon Personalize
- Amazon Polly
- Amazon Rekognition
- Amazon Textract
- Amazon Transcribe
- Amazon Translate
- AWS DeepLens
- AWS DeepRacer

## Analytics
- Athena
- EMR
- CloudSearch
- Elasticsearch Service
- Kinesis

## Application Integration
- Step Functions
- Amazon EventBridge
- Amazon MQ
- Simple Notification Service
- Simple Queue Service
- SWF

## AWS Cost Management
- AWS Cost Explorer
- AWS Budgets
- AWS Marketplace Subscriptions

## Customer Engagement
- Amazon Connect
- Pinpoint
- Simple Email Service

## Business Applications
- Alexa for Business
- Amazon Chime
- WorkMail

## Amazon RDS

- Dashboard
- Databases
- Query Editor
- Performance Insights
- Snapshots
- Automated backups
- Reserved instances
- Proxies

- Subnet groups
- Parameter groups
- Option groups
- Custom Availability Zones

- Events
- Event subscriptions

- Recommendations 20
- Certificate update

## Resources

Refresh

You are using the following Amazon RDS resources in the US East (N. Virginia) region (used/quota)

DB Instances (41/100)
- Allocated storage (0.13 TB/100 TB)
- Click here to increase DB instances limit

DB Clusters (25/40)

Reserved instances (0/100)

Snapshots (170)
- Manual (1/100)
- Automated (137)

Recent events (45)

Event subscriptions (0/20)

Parameter groups (22)
- Default (9)
- Custom (1[illegible]00)

Option groups (7)
- Default (7)
- Custom (0/20)

Subnet groups ([redacted])

Supported platforms VPC

[redacted]

## Create database

Amazon Relational Database Service (RDS) makes it easy to set up, operate, and scale a relational database in the cloud.

Restore from S3 | Create database

Note: your DB instances will launch in the US East (N. Virginia) region

## Service health

View service health dashboard

## Recommended for you

**Automatic Failover with RDS Multi-AZ**

Enable Multi-AZ configurations on your production workloads for a more resilient DR strategy ensuring zero data-loss and automated backups. Learn more

**RDS Performance Insights**

Quickly assess load on your database and take faster action with an easy-to-use performance dashboard. Learn more

**Fast Cross-Region Disaster Recovery**

Deploy your database across multiple AWS regions with no impact on performance, low lag and cross-region scaling of reads with Global Database. Learn more

**Faster Queries with Parallel Query**

Speed up queries without having to copy data into a separate system, while maintaining high throughput for your core transactional workload. Learn more

## Additional information

- Getting started with RDS
- Overview and features
- Documentation

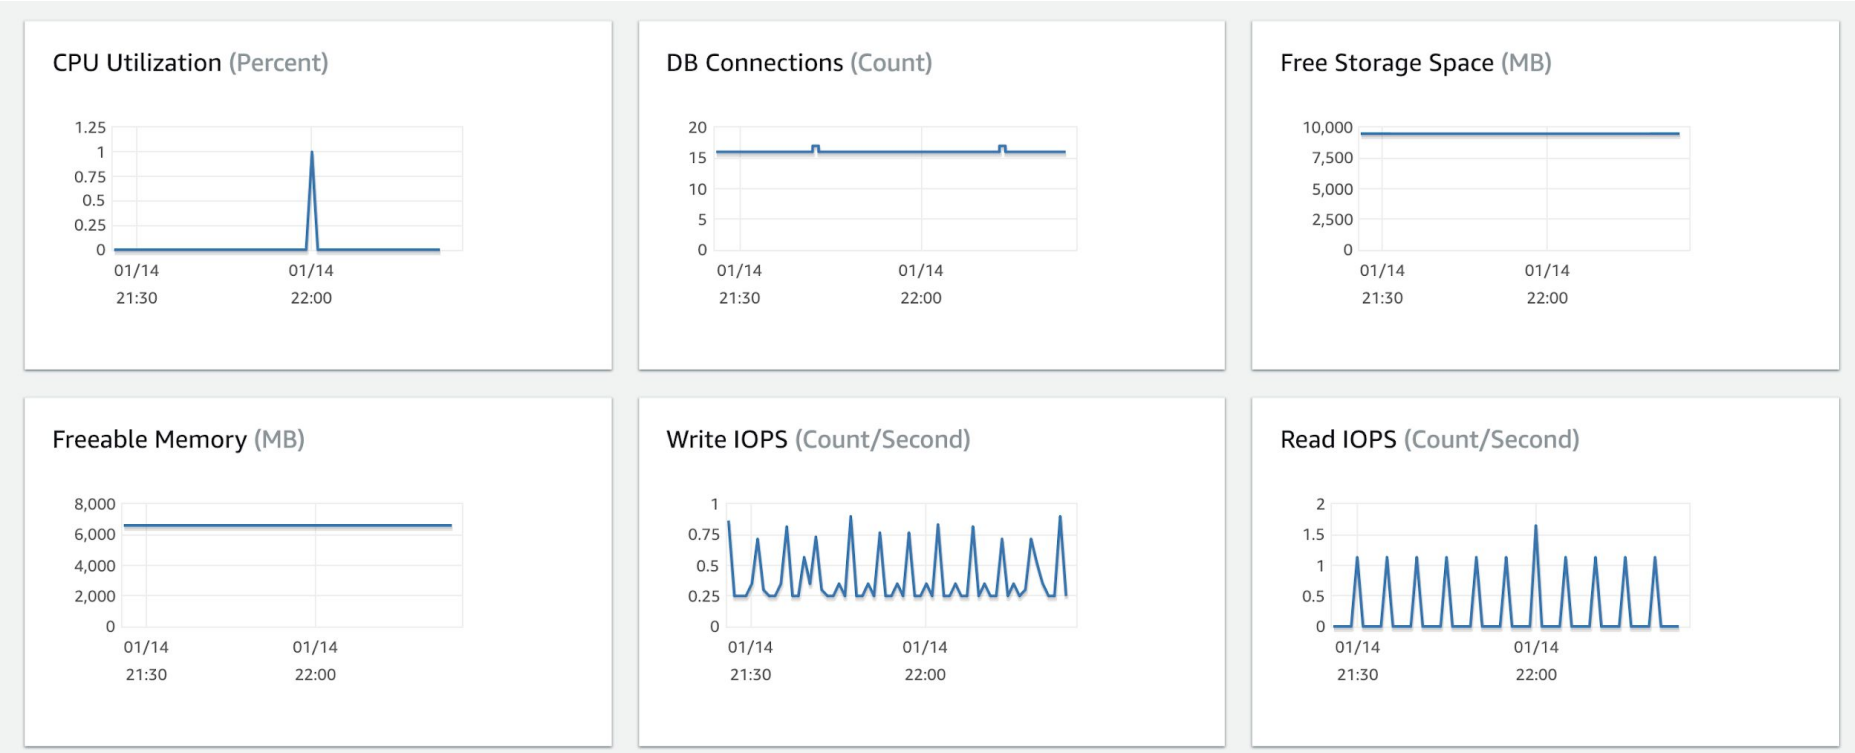
CPU Utilization (Percent)
1.25
1
0.75
0.5
0.25
0
01/14 21:30
01/14 22:00
DB Connections (Count)
20
15
10
5
0
01/14 21:30
01/14 22:00
Free Storage Space (MB)
10,000
7,500
5,000
2,500
0
01/14 21:30
01/14 22:00
Freeable Memory (MB)
8,000
6,000
4,000
2,000
0
01/14 21:30
01/14 22:00
Write IOPS (Count/Second)
1
0.75
0.5
0.25
0
01/14 21:30
01/14 22:00
Read IOPS (Count/Second)
2
1.5
1
0.5
0
01/14 21:30
01/14 22:00

# Datadog

# Datadog

Search... | Configure template variables

| Name | Tag or Attribute | Default Value | |
|---|---|---|---|
| environment | environment | staging | In 8/10 expressions Remove From All Add to All |

Add Variable + Done

Group | Notes & Links | Timeseries | Query Value 3.14 | Table | Top List | SLO Summary 99% BETA | Heat Map | Check Status 2 9 | Change | Host Map | Scatter Plot | Distribu

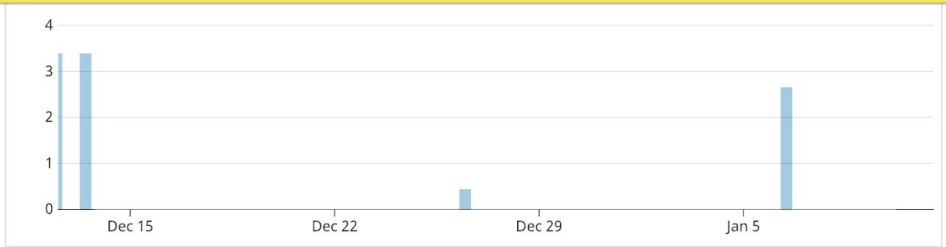

# Datadog

Edit Overview Correlations 1mo Past 1 Month

## 1 Select your visualization

Timeseries | Query Value | Table | Heat Map | Scatter Plot | Distribution | Top List | Change | Host Map

## 2 Graph your data

Edit JSON Share Graphing help

Metric system.cpu.user from (everywhere) avg by (everything) as... </>

Display: Lines Color: Classic Style: Solid Stroke: Normal Advanced...

Graph additional: Metrics | Log Events | Analyzed Spans | Live Processes | Network Traffic | RUM Events

- Event Overlays
- Markers
- Y-Axis Controls

## 3 Give your graph a title (or leave blank for suggested title)

Avg of system.cpu.user over *

Cancel Save

# Datadog

Edit Overview Correlations 1mo Past 1 Month

1 **Select your visualization**

Timeseries | Query Value | Table | Heat Map | Scatter Plot | Distribution | Top List | Change | Host Map

2 **Graph your data**

Edit JSON Share Graphing help

Metric system.cpu.user from (everywhere) avg by (everything) as... </>

Display: Lines Color: Classic Style: Solid Stroke: Normal Advanced...

Graph additional: Metrics | Log Events | Analyzed Spans | Live Processes | Network Traffic | RUM Events

- Event Overlays
- Markers
- Y-Axis Controls

3 **Give your graph a title** (or leave blank for suggested title)

Avg of system.cpu.user over *

## 2 Graph your data

Edit JSON Share

Metric

Display: Lines

Graph additional:

Events

Event Overlay

Markers

Y-Axis Contro

system.cpu.user

system

system.io.r_s

system.io.w_s

system.load.1

system.load.5

system.uptime

system.io.util

system.load.15

system.cpu.idle

system.cpu.user

system.io.await

system.io.rkb_s

## 3 Give your graph

Avg of system.cpu.u

## 2 Graph your data

Edit JSON Share

Metric

Display: Lines

Graph additional:

Event Overlay

Markers

Y-Axis Contro

system.cpu.user

kuber

kubernetes.cpu.limits

kubernetes.memory.rss

kubernetes.memory.swap

kubernetes.cpu.capacity

kubernetes.cpu.requests

kubernetes.memory.cache

kubernetes.memory.usage

kubernetes.pods.expired

kubernetes.pods.running

kubernetes.io.read_bytes

kubernetes.memory.limits

## 3 Give your graph

Avg of system.cpu.u

Kubernetes Namespaces

Edit Widgets +

2d The Past 2 Days

Search... | Save or select views | $namespace | $environment staging | $destination_workload_namespace *

## Total Pods Per Namespace

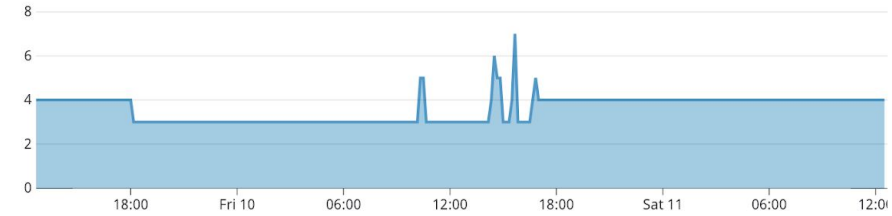


## Unavailable Pods

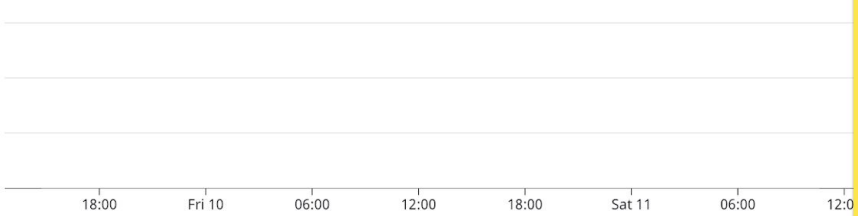


## CPU Usage Per Namespace

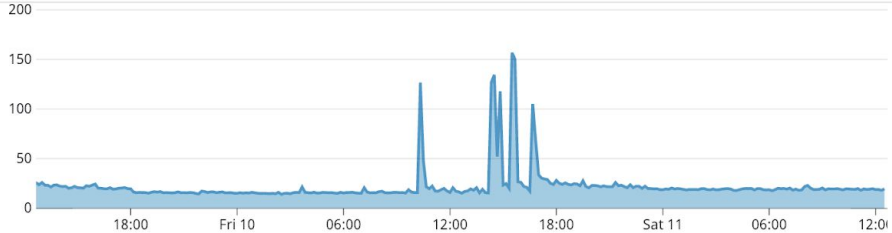


## Memory Usage Per Namespace

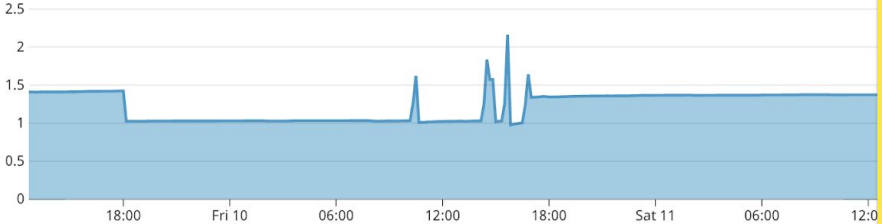


## Total CPU Requests Per Namespace


## Total Memory Requests Per Namespace

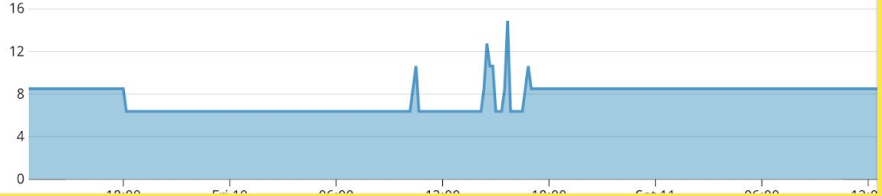

# Kibana

kibana

- Discover
- Visualize
- Dashboard
- Timelion
- Machine Learning
- Canvas
- Sentinl
- Infrastructure
- Logs
- APM
- Dev Tools
- Monitoring
- Management

D Default

Collapse

**960** hits

>_ Search... (e.g. status:200 AND e

Add a filter +

segment-test-framework*

Selected fields

? _source

Available fields

Popular

t @version

t _type

t kubernetes.labels.app

t log

t method

t response_code

t tags

@timestamp

t _id

t _index

# _score

t authority

t bytes_received

t bytes_sent

# Logs

```
LOGGER.error("Here’s the exact information you will need {}" , request);
```

# Kubectl Commands -

- `Kubectl staging get pods -n namepspace`
- `Kubectl staging describe pods -n namepspace`
- `Kubectl staging logs pod_id -n namepspace -c container -f`

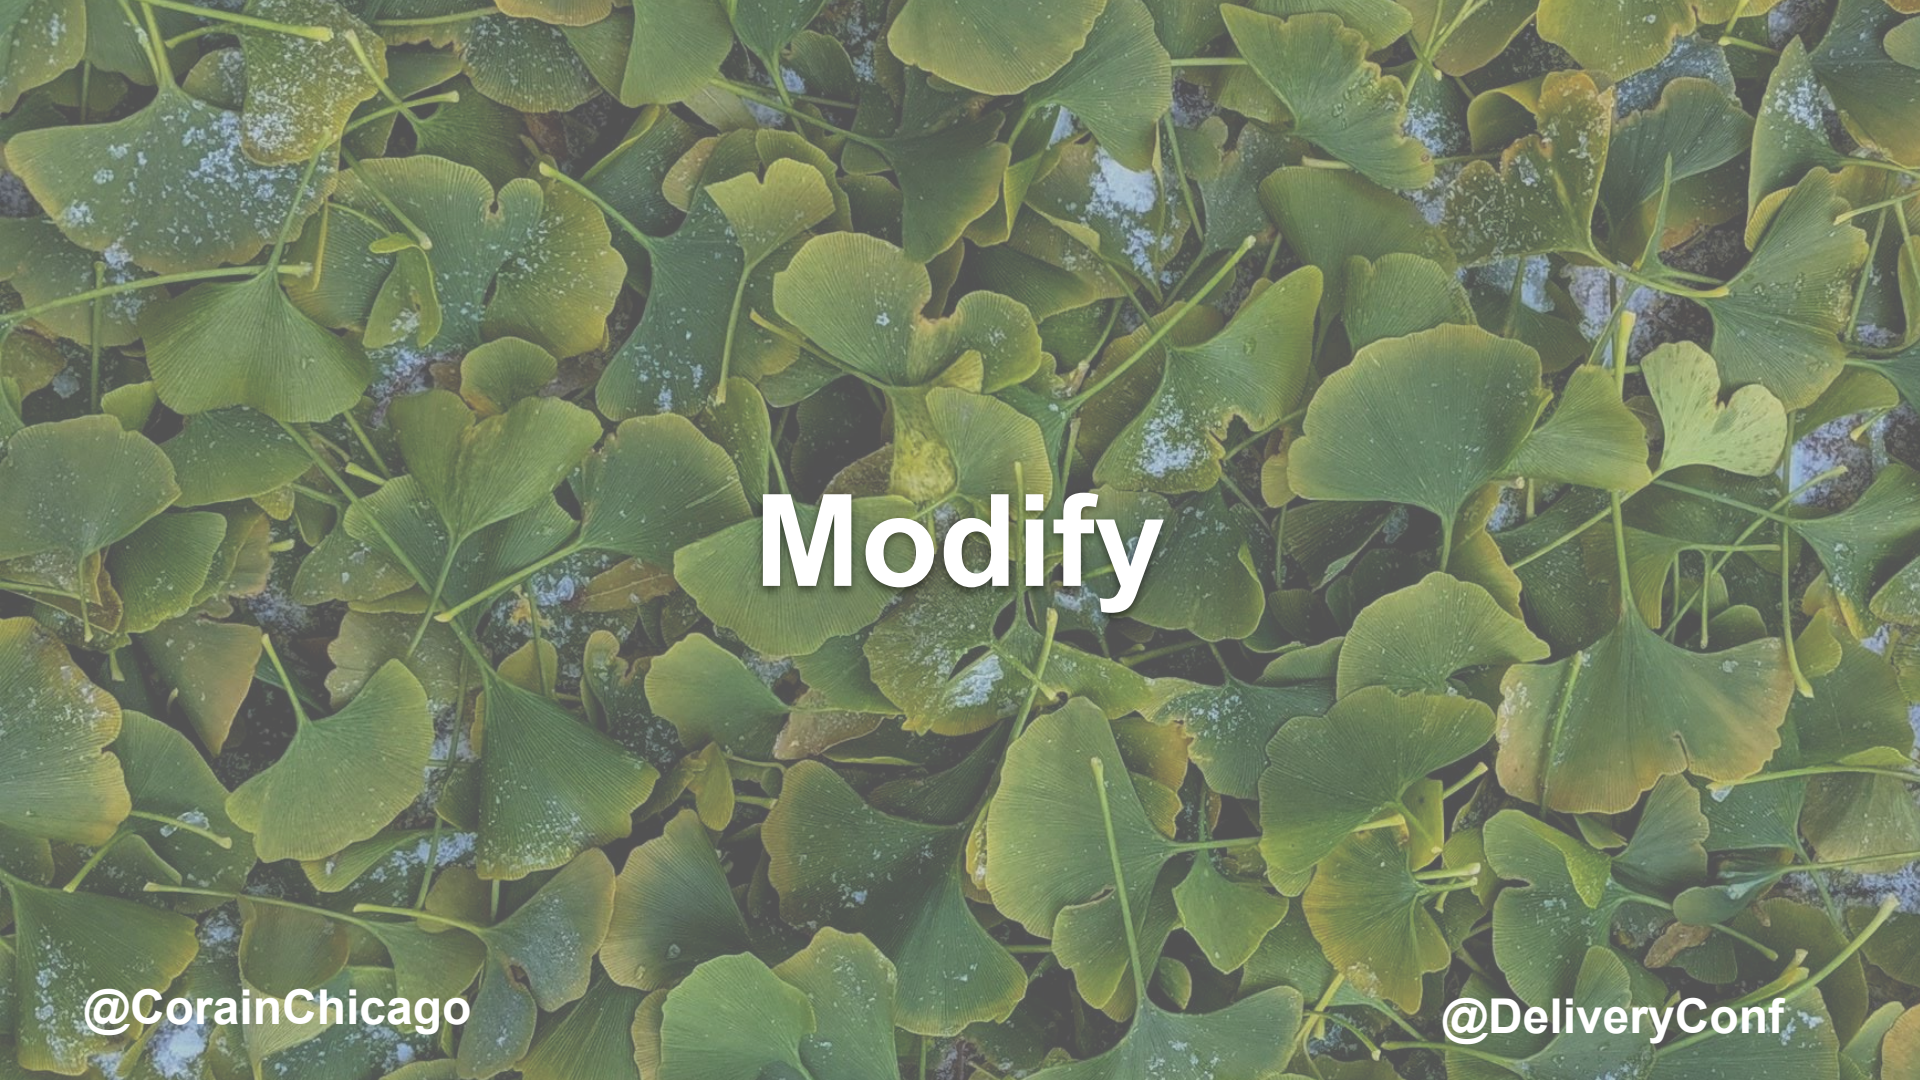
Modify
@CorainChicago
@DeliveryConf

# Repeat

1. ~~Tutorial~~
2. Code
3. PR
4. PR Comments - ~~SRE~~ Dev
5. PR update
6. ~~PR Comments - SRE~~
7. ~~PR update~~
8. ~~1:1 with SRE (briefer)~~
9. ~~Code~~
10. ~~PR~~
11. ~~PR Comments - SRE~~
12. ~~PR Update~~
13. ~~PR Comments - SRE~~
14. ~~PR Update~~
15. ~~1:1 Apply - SRE~~
16. SRE Ticket to apply

# Our story ends with…

(Also known as now)

# Thank you!

(We're hiring!)


ActiveCampaign

# Thank you!

**Cora Fedesna**

(she/her)

@CorainChicago

noti.st/corainchicago

@DeliveryConf