
foremast Documentation

Release 5.15.4

Foremast

Oct 25, 2022

Contents

1	Contents	1
1.1	Getting Started	1
1.2	Pipeline Flow and Examples	3
1.3	Configuration Files	5
1.4	Advance Usages	89
1.5	Infrastructure	91
1.6	Foremast - AWS Lambda Pipelines	96
1.7	Foremast - AWS S3 Pipelines	110
1.8	Deploy Spinnaker Using Halyard	111
1.9	How To Contribute	123
1.10	How To Create Releases	125
1.11	src	126
1.12	Changelog	135
2	Spinnaker Foremast	261
2.1	Why Foremast?	261
2.2	Foremast Features	262
2.3	Getting Started	262
2.4	More Details	262
3	Indices and tables	265

CHAPTER 1

Contents

1.1 Getting Started

- *Quick Start Guide*
 - *Installation*
 - * *Method 1 - Using pip (Preferred)*
 - * *Method 2 - Using git*
 - *Configuration Files*
 - *Running*

This getting started guide will walk through the process of using Foremast to create an application in Spinnaker and dynamically generate a basic Spinnaker pipeline.

Getting started with Foremast consists of the following steps:

1. Setting up configuration files
2. Installing Foremast
3. Setting up the variables
4. Running Foremast

1.1.1 Quick Start Guide

In this section, we will install, configure and run Foremast to create a basic pipeline.

Installation

Setting up the environment

```
$ pip3 install virtualenv
$ virtualenv -p $(which python3) venv
$ source venv/bin/activate
```

Method 1 - Using pip (Preferred)

```
$ pip install foremast
```

Method 2 - Using git

```
$ git clone https://github.com/foremast/foremast.git
$ cd foremast
$ pip install -U .
```

Configuration Files

Create a `runway` and `.foremast` directory and go into `runway` directory.

```
$ mkdir runway .foremast
```

Create `pipeline.json` in `runway` directory

```
{
    "deployment": "spinnaker",
    "env": [ "dev" ]
}
```

Create `application-master-dev.json` in `runway` directory.

```
{
    "app": {
        "instance_type": "t2.micro"
    },
    "asg": {
        "max_inst": 1,
        "min_inst": 1
    },
    "regions": [
        "us-east-1"
    ]
}
```

Go to `.foremast` directory and create the `foremast.cfg` file.

```
[base]
domain = example.com
envs = dev,prod
```

(continues on next page)

(continued from previous page)

```
regions = us-east-1
gate_api_url = http://gate.example.com:8084
```

You should now see something similar structure.

```
$ tree -a
.
└── .foremast
    └── foremast.cfg
    └── runway
        ├── application-master-dev.json
        └── pipeline.json
2 directories, 3 files
```

Running

Now from within the root directory, run `foremast-pipeline`.

```
$ GIT_REPO=hello PROJECT=world RUNWAY_DIR=runway/ foremast-pipeline
```

This will create an application in Spinnaker named `helloworld` along with a simple pipeline.

1.2 Pipeline Flow and Examples

- *Default Pipeline Flow*
- *Custom Pipelines*
- *Example Workflow*

Foremast generates a single pipeline per region. The pipeline is designed to allow deploying to multiple environment with checkpoints between each transition.

1.2.1 Default Pipeline Flow

The below flow can repeat for as many environments as defined in the configs. In general, most applications repeat a common set of stages 2-3 times. Typically, the same way to deploy to dev (if used), stage, and production.

1. Configuration
 - This stages defines the Jenkins trigger, property files, and pipeline-wide notifications
2. Bake
 - Bakes an AMI the specified AMI ID
3. Infrastructure Setup [\$env]
 - Calls a Jenkins job to run the `prepare-infrastructure` Foremast command against a specific account.
 - Sets up AWS infrastructure such as ELB, IAM roles, S3 bucket, and DNS needed for an application

4. Deploy \$env

- Uses Spinnaker to create a cluster and server group in specific account.
- The behavior of this stage is largely based on the [application-master-\\$account.json](#) configs.

5. Attach Scaling Policy [\$env]

- If a scaling policy is defined in [application-master-\\$account.json](#), attaches it to the deployed server group
- If no policy is defined, this stage is excluded

6. Checkpoint \$next-env

- A manual checkpoint stage. This requires human intervention to approve deployment to the next environment.

Stages 3-6 repeat for each environment/account defined in [pipeline.json](#).



The default generated pipeline should look like the above image. This is the basic *bake -> infrastructure -> deploy -> checkpoint* pipeline described above.

1.2.2 Custom Pipelines

You can specify an external templates directory in [foremast.cfg / config.py](#). Templates in an external directory will need to have the same directory structure and naming as the default templates. if `templates_path` is set in [foremast.cfg / config.py](#), Foremast will first see if the file exists there. If not, it will fall back to the provided templates.

If you need to add more stages or change the defaults, this is all possible via external templates. Please the [foremast-templates](#) repo for examples on the templates.

1.2.3 Example Workflow

Most Foremast users have a detailed workflow for using Foremast internally. Feel free to copy this example workflow or use it as inspiration for your own. You can view all of our internal templates on the [foremast-templates](#) repo.



1. The [application-master-\\$account.json](#) and [pipeline.json](#) are bundled directly with the application code
2. Developer makes a change to one of those configs and pushes to the application's git repository
3. A server-side git hook detects a change and triggers a Jenkins job to run Foremast `prepare-app-pipeline`
 - Regenerates the application and pipeline in Spinnaker
4. Build application artifacts using a Jenkins job and stored as an RPM
5. Spinnaker triggers detect a completed Jenkins job and starts a new deployment pipeline

1. Bake an AMI using built RPM
2. Jenkins Stage runs Foremast `prepare-infrastructure`
 - Creates the AWS ELB, SG, S3 bucket, and IAM roles
3. Jenkins Stage tags the source repository with AMI info
4. Deploy the generated AMI to desired environments
5. Jenkins Stage runs Quality Engineering tests against deployed application
6. Jenkins Stage tags the source repository with deployment information
7. Jenkins Stage attaches defined scaling policies
8. Manual judgment before continuing to the next environment

1.3 Configuration Files

1.3.1 foremast.cfg / config.py

- *Purpose*
- *Example Configuration*
- *Configuration Locations*
- *Specifying a Configuration file*
- *Configuration Details*
 - `[base]`
 - * `domain`
 - * `envs`
 - * `env_configs Keys`
 - * `aws_types`
 - * `gcp_types`
 - * `aws_manual_types`
 - * `gcp_manual_types`
 - * `regions`
 - * `ami_json_url`
 - * `gitlab_url`
 - * `gate_api_url`
 - * `templates_path`
 - * `runway_base_path`
 - * `default_run_as_user`
 - * `default_ec2_securitygroups`

```
* default_elb_securitygroups
* default_securitygroup_rules
* ec2_pipeline_types
* gate_client_cert
* gate_ca_bundle
- [credentials]
  * gitlab_token
  * slack_token
  * gate_authentication Keys
  * google_iap Keys
- [whitelists]
  * asg_whitelist
- [formats]
  * domain
  * app
  * dns_elb
  * s3_bucket
  * jenkins_job_name
- [task_timeouts]
  * default
  * envs
- [gcp]
  * envs
```

Purpose

This configuration holds information necessary for running foremast such as auth tokens, URLs, whitelists etc

Example Configuration

```
; foremast.cfg
[base]
domain = example.com
envs = dev,stage,prod
regions = us-east-1,us-west-2
ami_json_url = http://s3.bucketname.com/ami_lookup.json
git_url = https://git.example.com
gate_api_url = http://gate-api.example.com:8084
templates_path = ../../foremast-templates
default_run_as_user = trigger_runner
```

(continues on next page)

(continued from previous page)

```

default_securitygroup_rules = { "bastion" : [ { "start_port": "22", "end_port": "22",
    ↪"protocol": "tcp" } ], "serviceapp" : [ { "start_port": "8080", "end_port":
    ↪"8080", "protocol": "tcp" } ] }

[credentials]
gitlab_token = 123token23423343
slack_token = 123slack3203120312

[whitelists]
asg_whitelist = application1,application2

[formats]
app = {project}{repo}
dns_elb = lb-{project}{repo}.{env}.{domain}
s3_bucket = secret-{env}-{project}

[task_timeouts]
default = 120
envs = { "dev" : { "deleteScalingPolicy": 240} }

```

```

# config.py

CONFIG = {
    'base': {
        'domain': 'example.com',
        'envs': 'dev,stage,prod',
        'regions': 'us-east-1,us-west-2',
        'vpc_name': 'vpc',
        'ami_json_url': 'http://s3.bucketname.com/ami_lookup.json',
        'git_url': 'https://git.example.com',
        'gate_api_url': 'http://gate-api.example.com:8084',
        'templates_path': '../..//foremast-templates',
        'default_run_as_user': 'trigger_runner',
        'default_securitygroup_rules': {
            'bastion': [{}'start_port': '22', 'end_port': '22', 'protocol': 'tcp']],
            'serviceapp': [{}'start_port': '8080', 'end_port': '8080', 'protocol': 'tcp'],
        ↪' ],
        },
    },
    'credentials': {
        'gitlab_token': '123token23423343',
        'slack_token': '123slack3203120312',
        'gate_authentication': {
            'google_iap': {
                'enabled': False,
                'oauth_client_id': 'leeroyjenkins.apps.googleusercontent.com',
                'sa_credentials_path': '/tmp/service-account-creds.json'
            }
        }
    },
    'whitelists': {
        'asg_whitelist': 'application1,application2',
    },
    'formats': {

```

(continues on next page)

(continued from previous page)

```
'app': '{project}{repo}',
'dns_elb': 'lb-{project}{repo}.{env}.{domain}',
's3_bucket': 'secret-{env}-{project}',
},
'timeouts': {
    'default': 120,
    'envs': { 'dev': { 'deleteScalingPolicy': 240 } },
}
}
```

Configuration Locations

Foremast will look in the following locations, in order, for the `foremast.cfg` or `config.py` config file.

1. `./foremast/foremast.cfg`
2. `~/.foremast/.foremast.cfg`
3. `/etc/foremast/foremast/cfg`
4. `./config.py`

Specifying a Configuration file

Optionally, it is possible to specify the location of a configuration file for Foremast to use by setting the `FOREMAST_CONFIG_FILE` environment variable. This is useful if you do not store your config file in one of the locations listed above, or if you need to toggle between multiple configuration files to support different configurations.

Example: A config file for two different Spinnaker Instances

```
# Generate pipeline for spinnaker1
FOREMAST_CONFIG_FILE=config-spinnaker1.py
foremast generate pipeline
# Generate pipeline for spinnaker2
FOREMAST_CONFIG_FILE=config-spinnaker2.py
foremast generate pipeline
```

Configuration Details

[base]

Sections for base information such as urls and general configurations

domain

The base domain of your applications. Used for generating DNS

Required: Yes

envs

Comma delimited list of environments/applications that will be managed with Foremast for AWS. See section [\[gcp\]](#) for GCP Environments.

Example: dev, stage, prod

Required: Yes

env_configs Keys

Nested dictionary of environment names along with environment features

Type: Object

Default: None

Example Configuration:

```
{
    'env_configs': {
        "build": {
            "enable_approval_skip": True
        },
        "data": {
            "enable_approval_skip": False
        },
        "media": {
            "enable_approval_skip": False
        },
        "stage": {
            "enable_approval_skip": True
        },
        "prod": {
            "enable_approval_skip": False
        },
        "prodpp": {
            "enable_approval_skip": False
        }
    }
}

``enable_approval``
^^^^^^^^^^^^^^^^^^^^^

Determines if approval skips are allowed in this environment. Allows admins to_
→ ultimately enforce
deployment approvals in templates

| *Type*: boolean
| *Default*: ``False``
```

aws_types

Warning: *aws_types* replaced *types* beginning in Foremast 5.x when GCP support was added. It is recommended to migrate from the deprecated *types* configuration option to the new *aws_types*.

List of foremast managed Pipeline types to allow for AWS deployments

Type: str

Example: ec2, lambda

Default: ec2, lambda, s3, datapipeline, rolling

Required: No

gcp_types

List of foremast managed Pipeline types to allow for GCP deployments

Type: str

Example: cloudfunction

Default: cloudfunction

Required: No

aws_manual_types

Warning: *aws_manual_types* replaced *manual_types* beginning in Foremast 5.x when GCP support was added. It is recommended to migrate from the deprecated *manual_types* configuration option to the new *aws_manual_types*.

List of pipeline types that will trigger Foremast's manual pipeline template feature. When Foremast Infrastructure features are used the pipeline types listed here will create AWS infrastructure. See *advanced_manual_pipelines* for more details on this feature.

Type: str

Example: manual, custom_pipeline_name

Default: manual

Required: No

gcp_manual_types

List of pipeline types that will trigger Foremast's manual pipeline template feature. When Foremast Infrastructure features are used the pipeline types listed here will create GCP infrastructure. See *advanced_manual_pipelines* for more details on this feature.

Type: str

Example: gke, custom_pipeline_name

Default: “”

Required: No

regions

Comma delimited list of AWS regions managed by Foremast

Example: us-east-1,us-west-2

Required: Yes

ami_json_url

FQDN of where to query for AMI ID look ups. See [ami-lookup.json](#) for more details

Required: No

gitlab_url

FQDN of gitlab. Will be used for handling API calls to Gitlab

Required: No

gate_api_url

FQDN Of your spinnaker Gate instance. This is where all API calls to Spinnaker will go

Required: Yes

templates_path

Path to custom templates directory. If provided, Foremast will first look in this directory for any templates. This can be an absolute path, or a path relative to where you are running the Foremast commands. See [Pipeline Flow and Examples](#) for more details on custom templates.

Required: No

runway_base_path

Base path to use when looking for custom runway directories. If provided, Foremast will first look for Foremast runway files in this directory. This is useful if you have a different folder or location to store pipeline configuration values.

Type: str

Default: runway

Required: No

default_run_as_user

Default user to run pipelines as. This is needed for leveraging service accounts in Fiat.

Type: str

Default: null

Required: No

default_ec2_securitygroups

Comma separated list or json of EC2 security groups to include for all deployments. If a comma separated list is given, the groups are applied to all environments. If a json is provided, it assigns groups only to the specified environment.

Required: No

Example: office,test_sg,example

Example (json): {"dev": ["sg1", "sg2"], "stage": ["sg3"]}

default_elb_securitygroups

Comma separated list or json of ELB security groups to include for all deployments. If a comma separated list is given, the groups are applied to all environments. If a json is provided, it assigns groups only to the specified environment.

Required: No

Example: test_sg,example_elb_sg

Example (json): {"dev": ["sg1", "sg2"], "stage": ["sg3"]}

default_securitygroup_rules

Security group rules that should be included by default for the application specific group. If `$self` is used as the security group name, it will self-reference to its own application name.

Required: No

Example: { "bastion" : [{ "start_port": "22", "end_port": "22", "protocol": "tcp" }] }

ec2_pipeline_types

`foremast.consts.EC2_PIPELINE_TYPES = ('ec2', 'rolling')`

Comma separated list of Pipeline Types to treat as EC2 deployments.

This is useful when defining custom Pipeline Types. When Pipeline Type matches, EC2 specific data is used in deployment, such as Auto Scaling Groups and Availability Zones.

Default: ec2,rolling

Required: No

Example: ec2,infrastructure,propeller

gate_client_cert

If accessing Gate via x509 certificate authentication, this value provides the local path to the certificate. Only PEM certs are supported at this time (containing both the key and certificate in the PEM).

Required: No

Example: /var/certs/gate-cert.pem

gate_ca_bundle

If accessing Gate via x509 leveraging a custom certificate authority (such as acting as your own CA), this value provides the local path to the CA bundle. It is recommended to use an existing CA Bundle and append your CA certificate to it (<https://certifi.io/en/latest/>)

Required: No

Example: /var/certs/CA.pem

[credentials]

Section for handling credential configurations such as tokens, usernames, and passwords

gitlab_token

Gitlab token used for authentication in Foremast

Required: No

slack_token

Slack token used for authentication when sending Slack messages from Foremast

Required: No

gate_authentication Keys

Credential Provider Object used to authenticate to Gate

Type: Object

Default: None

Example Configuration:

```
{
    'credentials': {
        'gate_authentication': {
            'google_iap': {
                'enabled': False,
                'oauth_client_id': 'some_id.apps.googleusercontent.com',
                'sa_credentials_path': '/tmp/google-service-account.json'
            },
            'github': {
                'token': '<GITHUB_TOKEN>'
            }
        }
    }
}
```

google_iap Keys

We currently support in addition to x509, Google Identity Aware Proxy authentication.

Determines if this authentication method should be used or not.

Type: boolean

Default: False

oauth_client_id

Application Client ID using Identity Aware Proxy. Can be found in the Google Cloud Console

Type: string

Default: None

`sa_credentials_path`

Path to Google Cloud Service Account used to Authenticate to Identity Aware Proxy. Must be added to IAP in GCP console to grant permission.

Type: string

Default: None

[whitelists]

Sections for configuring whitelist info

`asg_whitelist`

Comma delimited list of applications to whitelist from ASG rules

Required: No

[formats]

Section handling the naming convention of applications, elb, iam, s3 buckets and other services.

The most common sections are shown. The complete list of sections and defaults are defined by the underlying library [foremast-utils](#).

Any of the possible variables below can be used as the value.

- domain organization domain
- env dev, qa, production, etc
- project lowercase git group/organization
- repo lowercase git project/repository
- raw_project git group/organization
- raw_repo git project/repository

`domain`

A string of your organization's domain

Default: example.com

Required: No

app

A string of the format of your application

Default: {repo}{project}

Required: No

dns_elb

An FQDN of your application's Elastic Load Balancer (ELB)

Default: {repo}.{project}.{env}.{domain}

Required: No

s3_bucket

An string of your base S3 bucket name

Default: archaius-{env}

Required: No

jenkins_job_name

An string of the format of the application's jenkins job name

Default: {project}_{repo}

Required: No

[task_timeouts]

Section handling customization of task timeouts when communicating with Spinnaker. Timeouts can vary per environment and per task.

default

The default task timeout value

Default: 120

Required: No

envs

A json object keyed by environment name. Each value should be a json object keyed by task name. This section only applies to AWS environments.

Default: 120

Required: No

[gcp]

Section handling GCP infrastructure and authentication configuration options.

envs

A json object keyed by environment name. Each value should be a json object that defines the GCP environment's structure. The property *service_account_project* defines which project is used by Foremast when creating service accounts. You should use different a *service_account_project* for each environment to ensure IAM permissions are not granted between environments. See the page [GCP Credentials](#) for more info on setting up GCP credentials for Foremast.

Default: None

Required: Yes

Example structure:

```
{  
    'stage': {  
        'organization': 'your-org.com',  
        'service_account_project': 'project-id-for-creating-service-accounts-stage',  
        'service_account_path': '/path/to/service/account/used/by/foremast-stage.json'  
    },  
    'prod': {  
        'organization': 'your-org.com',  
        'service_account_project': 'project-id-for-creating-service-accounts-prod',  
        'service_account_path': '/path/to/service/account/used/by/foremast-prod.json'  
    },  
}
```

1.3.2 pipeline.json

- *Purpose*
- *Example Configuration*
- *Configuration Details*
 - *type*
 - *owner_email*
 - *documentation*
 - *notifications Block*
 - * *email*
 - * *slack*
 - *pipeline_notifications Array*
 - * *pipeline_notifications*
 - *promote_restrict*
 - *base*

```
- env
- image Block
  * bake_instance_type
  * root_volume_size
- lambda Block
  * app_description
  * handler
  * runtime
  * vpc_enabled
- services Block
  * athena
  * cloudformation
  * cloudwatchlogs
  * cloudwatch
  * datapipeline
  * dynamodb
  * elasticache
  * elasticsearch
  * emr
  * firehose
  * glue
  * kinesis
  * kms
  * lambda
  * mediaconvert
  * parameterstore
  * rds-db
  * rds-data
  * redshift-data
  * s3
  * sdb
  * secretsmanager
  * ses
  * sns
  * sqs
  * xray
```

```
* gcp_roles
- chaos_monkey Block
    * enabled
    * mean_time
    * minimum_time
    * exceptions
- instance_links Block
- permissions Block
    * read_roles
    * write_roles
- traffic_guards Block
    * accounts
- cloudfunction Block
    * project_name
    * entry_point
    * runtime
```

Purpose

This configuration file is used for defining pipeline settings that affect the pipeline as a whole, not a specific account/environment.

Example Configuration

```
{
    "type": "ec2",
    "owner_email": "",
    "documentation": "",
    "notifications": {
        "email": "",
        "slack": ""
    },
    "pipeline_notifications": [],
    "promote_restrict": "none",
    "base": "tomcat8",
    "env": ["stage", "prod"],
    "primary_region": "us-east-1",
    "image": {
        "bake_instance_type": "t2.small",
        "root_volume_size": 6,
        "builder": "ebs"
    },
    "lambda": {
        "app_description": "default description",
        "runtime": "java8",
```

(continues on next page)

(continued from previous page)

```

    "handler": "main",
    "vpc_enabled": false,
    "package_type": "zip"
},
"pipeline_files": [],
"chaos_monkey": {
    "enabled": false,
    "mean_time": 5,
    "minimum_time": 3,
    "exceptions": []
},
"instance_links": {},
"permissions": {
    "read_roles": [],
    "write_roles": []
},
"traffic_guards": {
    "accounts": []
},
"cloudfunction": {
    "project_name": "my-project*",
    "entry_point": "hello_get",
    "runtime": "python37"
}
}

```

Configuration Details

type

Specifies what type of pipeline to use for the application.

Type: string

Default: "ec2"

Options:

- "ec2" - Sets up an AWS EC2 pipeline and infrastructure
- "datapipeline" - Sets up an AWS Data Pipeline infrastructure
- "lambda" - Sets up an AWS Lambda pipeline and infrastructure
- "stepfunction" - Sets up an AWS Step Function pipeline and infrastructure
- "cloudfunction" - Sets up a GCP Cloud Function pipeline, infrastructure and deploys code
- "s3" - Sets up an AWS S3 pipeline and infrastructure
- "rolling" - Sets up a “rolling” style pipeline. Requires custom templates.
- "manual" - Sets up pipelines from raw Spinnaker Pipeline JSON; more info: [Configuration Files Advanced Usages](#).

owner_email

The application owners email address. This is not used directly in the pipeline but can be consumed by other tools

Type: string
Default: null

documentation

Link to the applications documentation. This is not used directly in the pipeline but can be consumed by other tools

Type: string
Default: null

notifications Block

Warning: notifications is deprecated, see "pipeline_notifications" instead

Where to send pipeline failure notifications

email

Email address to send pipeline failures (email must be configured in Spinnaker Echo)

Type: string
Default: null

slack

Slack channel to send pipeline failures (Slack must be configured in Spinnaker Echo)

Type: string
Default: null

pipeline_notifications Array

Where to send pipeline notifications. Notifications can be sent on several events including pipelines starting, completing and failing. Any supported notification option in Spinnaker can be defined, including Slack, Microsoft Teams, Bearychat, PubSub, Google Chat and Email.

pipeline_notifications

Array of notification definitions

Type: array
Default: []

Example Microsoft Teams:

Example Slack:

Example Email:

```
[  
  {  
    "level": "pipeline",  
    "type": "email",  
    "address": "jane.doe@who.com",  
    "cc": "jon.doe@optional.com",  
    "when": [  
      "pipeline.failed",  
      "pipeline.complete",  
      "pipeline.starting"  
    ]  
  }  
]
```

Example Google Cloud Pub/Sub:

```
[  
  {  
    "level": "pipeline",  
    "type": "pubsub",  
    "publisherName": "my-publisher",  
    "when": [  
      "pipeline.starting",  
      "pipeline.complete",  
      "pipeline.failed"  
    ]  
  }  
]
```

Example Google Chat:

```
[  
  {  
    "level": "pipeline",  
    "type": "googlechat",  
    "address": "https://chat.google.com/v1/spaces/my-google-chat-space",  
    "when": [  
      "pipeline.starting",  
      "pipeline.complete",  
      "pipeline.failed"  
    ]  
  }  
]
```

Example custom messages:

Some notification types support custom messages, which can be defined using the `messages` field:

```
[  
  {  
    /* First define your notification, e.g. slack or teams */  
    /* ... */  
    "message": {  
      "pipeline.complete": {  
        "text": "A pipeline finished, wow!"  
      },  
      "pipeline.failed": {  
        "text": "A pipeline has failed :("  
      },  
      "pipeline.starting": {  
        "text": "A pipeline started!"  
      }  
    }  
}
```

promote_restrict

Restriction setting for promotions to prod* accounts.

Type: string

Default: "none"

Options:

- "masters-only" - only masters/owners on a repository can approve deployments
- "members-only" - Any member of a repository can approve deployments
- "none" - No restrictions

base

The base AMI to use for baking the application. This can be an alias defined in `ami-lookup.json` or an AMI Id.

Type: string

Default: "tomcat8"

env

List of accounts that the application will be deployed to. Order matters as it defines the order of the pipeline. The accounts should be named the same as you have them in Spinnaker Clouddriver

Type: array

Default: ["stage", "prod"]

image Block

Holds settings for the baked image

bake_instance_type

Defines the instance type for Rosco (bake step) to use. This could help with issues of large and complex bakes. Refer to: <https://aws.amazon.com/ec2/instance-types/>

Type: string

Default: "t2.small"

root_volume_size

Defines the root volume size of the resulting AMI in GB

Type: number

Units: Gigabyte

Default: 6

lambda Block

Holds settings related to lambda deployments

app_description

Lambda function description

Type: string

Default: "default description"

handler

The function that Lambda calls to begin execution

Type: string

Default: "main"

runtime

The runtime environment for the Lambda function. Since value is passed directly to the lambda API new runtimes are automatically supported as they are released

Type: string

Default: "java8"

Options:

- "java8"

- "nodejs"

- "nodejs4.3"
- "nodejs6.10"
- "nodejs8.10"
- "python2.7"
- "python3.6"
- "dotnetcore1.0"
- "dotnetcore2.0"
- "nodejs4.3-edge"
- "go1.x"

vpc_enabled

Whether or not the Lambda function should use a VPC

Type: boolean

Default: false

services Block

Access to different Cloud Services will be added to an inline Policy for an IAM Role. Keys must match with a corresponding template in `src/foremast/templates/infrastructure/iam/key.json.j2`.

athena

Add Athena Query access.

Type: boolean

Default: false

cloudformation

Add CloudFormation access.

Type: boolean

Default: false

cloudwatchlogs

Add CloudWatch Logs access. Lambda Functions will automatically have this added.

Type: boolean

Default: false

cloudwatch

Add CloudWatch Limited access.

Type: boolean

Default: false

datapipeline

Allows a data pipeline to be trigger e.g. via lambda.

Type: boolean

Default: false

dynamodb

Add DynamoDB access to tables listed.

Type: array

Default: []

elasticache

Add Elasticache access to clusters listed.

Type: array

Default: []

elasticsearch

Add ElasticSearch access to domains listed.

Type: array

Default: []

emr

Add EMR Full access.

Type: boolean

Default: false

firehose

Add Firehose access to streams listed.

Type: array

Default: []

glue

Add Glue GetTable and GetDatabase access.

Type: boolean

Default: false

kinesis

Add Kinesis Streams access to streams listed.

Type: array

Default: []

kms

Add KMS Decrypt access to KMS keys listed.

Type: array

Default: []

lambda

Add Lambda access.

Type: boolean

Default: false

mediaconvert

Add MediaConvert Full access.

Type: boolean

Default: false

parameterstore

Add SSM ParameterStore PutParameter and GetParametersByPath access based on app name.

Type: boolean

Default: false

rds-db

Add RDS-DB Connect access to RDS DB Resources. Expects RDS DB User to match Spinnaker appname or use of Secrets Manager credentials for DB to connect. (<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>)

Type: array

Default: []

Example: ["db-12ABC34DEFG5HIJ6KLMNOP78QR", "*"]`

rds-data

Add RDS-Data APIs. By using the Data API for Aurora Serverless, you can work with a web-services interface to your Aurora Serverless DB cluster. The Data API doesn't require a persistent connection to the DB cluster. Instead, it provides a secure HTTP endpoint and integration with AWS SDKs. You can use the endpoint to run SQL statements without managing connections.

Requires AWS Secret Manager to be passed.

(<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html>)

Type: boolean

Default: false

redshift-data

Add Redshift-Data API. You can access your Amazon Redshift database using the built-in Amazon Redshift Data API. Using this API, you can access Amazon Redshift data with web services-based applications, including AWS Lambda, AWS AppSync, Amazon SageMaker notebooks, and AWS Cloud9.

The Data API doesn't require a persistent connection to the cluster. Instead, it provides a secure HTTP endpoint and integration with AWS SDKs. You can use the endpoint to run SQL statements without managing connections. Calls to the Data API are asynchronous.

The Data API uses either credentials stored in AWS Secrets Manager

(<https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>)

Type: boolean

Default: false

s3

Add S3 access to the provided Bucket. You may need to override default templates, see *templates_path*. To access other S3 Buckets, provide a list of names to add.

Type: boolean XOR array

Default: false

Example boolean:

```
{  
  "s3": true  
}
```

Example array:

```
{  
  "s3": ["other_bucket"]  
}
```

sdb

Add SimpleDB access to SimpleDB Domains listed.

Type: array

Default: []

secretsmanager

Add Secrets Manager access to the secrets listed.

Type: array

Default: []

ses

Add SES access.

Type: boolean

Default: false

sns

Add SNS access.

Type: boolean

Default: false

sqs

Add SQS access.

Type: boolean

Default: false

xray

Add AWS XRay access.

Type: boolean

Default: false

gcp_roles

Adds [GCP Roles](#) to the given projects.

Wildcards (*) are supported in the project_name field. For example *project-one** may match to *project-one-prod* or *project-one-stage* depending on what environment is being deployed to.

Type: array of objects

Default: None

Example:

```
"gcp_roles": [
    {
        "project_name": "project-one*",
        "roles": [
            "roles/secretmanager.secretAccessor",
            "roles/pubsub.subscriber"
        ]
    },
    {
        "project_name": "project-two*",
        "roles": [
            "roles/storage.objectViewer"
        ]
    }
]
```

chaos_monkey Block

Key that configures Chaos Monkey

enabled

Enable or disable Chaos Monkey

Type: boolean

Default: false

mean_time

Mean time between terminations. If mean_time is n, then the probability of a termination on each day is 1/n

Type: number

Default: 5

Units: Days

minimum_time

Minimum time between terminations

Type: number

Default: 3

Units: Days

exceptions

Accounts that Chaos Monkey will not affect

Type: array

Default: []

instance_links Block

Adds custom instance links to spinnaker. This takes a dictionary where the key is the name of the link and the value is the destination.

Example:

```
{  
    "instance_links": {  
        "health": ":8080/health",  
        "documentation": "http://example.com"  
    }  
}
```

permissions Block

Key that configures permissions for an application (leverages Fiat Roles/Groups) For more info, visit: <https://www.spinnaker.io/setup/security/authorization/>

read_roles

Roles that should have read permission to this application in Spinnaker

Type: array

Default: []

write_roles

Roles that should have write permission to this application in Spinnaker

Type: array

Default: []

traffic_guards Block

Key that configures Traffic Guards for an application

accounts

Accounts that Traffic Guards will be enabled for. Traffic Guards allow you to specify critical clusters that should always have active instances. If a user or process tries to delete, disable, or resize the server group, Spinnaker will verify the operation will not leave the cluster with no active instances, and fail the operation if it would.

Type: array

Default: []

cloudfunction Block

Holds settings related to GCP Cloud Function deployments

project_name

The project name. Wildcards are supported to ensure the correct project is used in each GCP Environment.

For example `my-project*` may match to `my-project-prod` and `my-project-stage` depending on the environment being deployed to.

Type: string

Default: None

Required: Yes

entry_point

The entry point of your code. Typically this is a function or method name.

Type: string

Default: None

Required: Yes

Example: `my_function`

runtime

The runtime your function is using. See the GCP docs for a full list of options.

Type: string

Default: None

Required: Yes

Example: `python37`

1.3.3 application-master-\$account.json

- *Purpose*
- *Example Configuration*
- *Configuration Details*
 - *app Block*
 - * *app_description*
 - * *approval_skip*
 - * *approval_timeout*
 - * *archaius_enabled*
 - * *custom_tags*

```
    · custom_tags Example
* eureka_enabled
* instance_profile
* instance_type
* lambda_concurrency_limit
* lambda_destinations
    · lambda_destinations Example
* lambda_dlq
    · lambda_dlq Example
* lambda_environment
    · lambda_environment Keys
    · lambda_environment Example
* lambda_filesystems
    · lambda_filesystems Example
* lambda_layers
* lambda_memory
* lambda_provisioned_throughput
* lambda_role
* lambda_subnet_count
* lambda_subnet_purpose
* lambda_timeout
* lambda_tracing
    · lambda_tracing Example
* cloudfunction_timeout
* cloudfunction_memory_mb
* cloudfunction_environment
* cloudfunction_allow_unauthenticated
* cloudfunction_iam_bindings
* cloudfunction_max_instances
* cloudfunction_ingress_type
* cloudfunction_vpc
    · connector
    · egress_type
* cloudfunction_event_trigger
    · event_type
```

```

        · resource

    - asg Block
        * hc_type
        * app_grace_period
        * max_inst
        * min_inst
        * ssh_keypair
        * subnet_purpose
        * enable_public_ips
        * scaling_policy
        * custom_scaling_policies
        * scheduled_actions

    - elb Block
        * access_log
            · access_log Keys
        * connection_draining_timeout
        * certificate
        * health
            · health Keys
        * idle_timeout
        * ports
            · ports Keys
            · ports Example
        * subnet_purpose
        * target

    - regions Key
    - deploy_strategy Key
    - security_group Block
        * description
        * elb_extras
        * instance_extras
        * ingress
        * egress
        * security_group Example

    - dns Block
        * ttl

```

```
- lambda_triggers
- datapipeline Block
  * name
  * description
  * activate_on_deploy
  * json_definition
- s3 Block
  * path
  * bucket_acl
  * bucket_name
  * bucket_policy
  * content_metadata
  * shared_bucket_master
  * shared_bucket_target
  * cors
    · cors Keys
  * encryption
    · encryption Keys
  * lifecycle
    · lifecycle Keys
  * logging
    · logging Keys
  * notification
    · notification Keys
  * tagging
    · tagging Keys
  * versioning
    · versioning Keys
  * website
    · website Keys
- stepfunction Block
  * statemachine_type
  * json_definition
  * tracing
  * logging_configuration
```

- *qe Block*
 - * *failure_action*
 - * *ignore_unstable_results*
 - * *jenkins_master*
 - * *test_job*
 - * *test_params*
 - * *stage_timeout*
 - * *wait_for_completion*
- *Example Quality Stage JSON Configuration*

Purpose

This configuration file holds infrastructure information for \$account. Each AWS account in your pipeline would need a separate application-master-\$account.json file. If your account is named dev, you would want an application-master-dev.json file.

Example Configuration

```
{
  "app": {
    "app_description": null,
    "approval_skip": false,
    "approval_timeout": null,
    "archaius_enabled": false,
    "canary": false,
    "custom_tags": {},
    "email": null,
    "eureka_enabled": false,
    "instance_profile": "{{ profile }}",
    "instance_type": "t2.micro",
    "lambda_concurrency_limit": null,
    "lambda_destinations": {},
    "lambda_dlq": {},
    "lambda_environment": {},
    "lambda_filesystems": [],
    "lambda_layers": [],
    "lambda_memory": "128",
    "lambda_role": null,
    "lambda_provisioned_throughput": null,
    "lambda_subnet_count": null,
    "lambda_subnet_purpose": "internal",
    "lambda_timeout": "30",
    "lambda_tracing": {},
    "cloudfunction_timeout": "60",
    "cloudfunction_memory_mb": "128",
    "cloudfunction_max_instances": "0",
    "cloudfunction_vpc_connector": null
  },
  "asg": {
    ...
  }
}
```

(continues on next page)

(continued from previous page)

```
"hc_type": "ELB",
"hc_grace_period": 180,
"app_grace_period": 0,
"max_inst": 3,
"min_inst": 1,
"ssh_keypair": null,
"subnet_purpose": "internal",
"enable_public_ips": null,
"provider_healthcheck": {
    "amazon": false
},
"scaling_policy": {},
"custom_scaling_policies": [],
"scheduled_actions": []
},
"elb": {
    "certificate": null,
    "policies": [],
    "listener_policies": [],
    "backend_policies": [],
    "idle_timeout": null,
    "access_log": {},
    "connection_draining_timeout": null,
    "health": {
        "interval": 20,
        "threshold": 2,
        "timeout": 10,
        "unhealthy_threshold": 5
    },
    "i_port": 8080,
    "i_proto": "HTTP",
    "lb_port": 80,
    "lb_proto": "HTTP",
    "subnet_purpose": "internal",
    "target": "TCP:8080"
},
"qe": {
},
"regions": {
    "us-east-1": {}
},
"deploy_strategy": "highlander",
"security_group": {
    "description": "Auto-Gen SG for {{ app }}",
    "egress": "0.0.0.0/0",
    "elb_extras": [],
    "ingress": {
    },
    "instance_extras": []
},
"dns": {
    "ttl": 60,
    "failover_dns": true,
    "region_specific": true
},
"lambda_triggers": [],
"s3": {
```

(continues on next page)

(continued from previous page)

```

"shared_bucket_master": false,
"bucket_name": "",
"path": "/",
"bucket_acl": "private",
"bucket_policy": {},
"content_metadata": [],
"cors": {
    "enabled": false,
    "cors_rules": [
        {
            "cors_headers": [],
            "cors_methods": [],
            "cors_origins": [],
            "cors_expose_headers": [],
            "cors_max_age": 600
        }
    ],
    "encryption": {
        "enabled": false,
        "encryption_rules": [
            {
                "ApplyServerSideEncryptionByDefault": {
                    "SSEAlgorithm": "AES256"
                }
            }
        ]
    },
    "lifecycle": {
        "enabled": false,
        "lifecycle_rules": [{}]
    },
    "logging": {
        "enabled": false,
        "logging_grants": [],
        "logging_bucket": "",
        "logging_bucket_prefix": "{{ app }}/"
    },
    "notification": {
        "enabled": false,
        "topic_configurations": [{}],
        "queue_configurations": [{}],
        "lambda_configurations": [{}]
    },
    "tagging": {
        "tags": {}
    },
    "versioning": {
        "enabled": false,
        "mfa_delete": "Disabled"
    },
    "website": {
        "enabled": false,
        "index_suffix": "index.html",
        "error_document": "404.html"
    }
},
"datapipeline": {
    "description": "",
    "activate_on_deploy": false,
    "json_definition": {}
}

```

(continues on next page)

(continued from previous page)

```
        },
        "stepfunction": {
            "tracing": {
                "enabled": false
            },
            "logging_configuration": {},
            "statemachine_type": "STANDARD",
            "json_definition": {}
        }
    }
```

Configuration Details

app Block

Top level key that contains information on the application and EC2 details

app_description

Describes the application.

Type: string
Default: null

approval_skip

Enable the ability to skip approval stage for a given environment. Must be enabled in foremast configs per environment to allow overrides.

Type: boolean
Default: false

approval_timeout

Enable the ability to override Spinnaker's default Stage Timeout (typically 72-hours) with a custom timeout specified in milliseconds. This is helpful to maintain cleaner pipelines, and fail pipelines not ready for the next environment. For example, 2 hours is represented as 7200000.

Type: int
Format: ms
Default: null

archaius_enabled

Setting this value to true will autoreate archaius pathing in a specified archaius S3 bucket.

Type: boolean
Default: false

custom_tags

Custom Tags to be used during deployment stages on resources such as ELBs and EC2s.

custom_tags Example

```
{
    "app": {
        "custom_tags": {
            "example_key": "example_value",
            "app_name": "application_name"
        }
    }
}
```

eureka_enabled

Setting this value to `true` will not create an ELB, DNS record, and set the ASG health check to EC2.

Type: boolean

Default: `false`

instance_profile

The instance profile to start EC2 instances with. Foremast creates default instance profile based on the default string. Specifying a different profile name assumes the profile exists.

Type: string

Default: `"${group}_${app}_profile"`

instance_type

The size/type of the EC2 instance. Uses Standard AWS instance names. See <https://aws.amazon.com/ec2-instance-types/> for details

Type: string

Default: `"t2.micro"`

lambda_concurrency_limit

Each region in your AWS account has a Lambda concurrency limit. The concurrency limit determines how many function invocations can run simultaneously in one region. The limit applies to all functions in the same region and is set to 1000 by default.

If you exceed a concurrency limit, Lambda starts throttling the offending functions by rejecting requests. Depending on the invocation type, you'll run into the following situations:

More info on limits can be found here: <https://docs.aws.amazon.com/lambda/latest/dg/limits.html>

lambda_destinations

This feature provides the ability to control what happens when a function is successful or fails e.g. if a specific function fails you may want to invoke another lambda function to perform some error management. In the past you would have to add this bespoke functionality into your code.

Destinations currently support following:

- * ARN of Lambda Function
- * ARN of SQS Queue
- * ARN of SNS Topic
- * ARN of Amazon EventBridge event bus

You may either an individual destination path OR one for both success and failure.

More details on lambda destinations can be found here: <https://aws.amazon.com/blogs/compute/introducing-aws-lambda-destinations/>

Type: Object

Default: {}

lambda_destinations Example

```
"lambda_destinations": {  
    "OnSuccess": { "Destination": "arn" },  
    "OnFailure": { "Destination": "arn" }  
}
```

lambda_dlq

A dead letter queue configuration that specifies the queue or topic where Lambda sends asynchronous events when they fail processing

Dead Letter Queues are supported in either SNS or SQS and pass in the ARN. See <https://docs.aws.amazon.com/lambda/latest/dg/invocation-async.html> for more details

Type: Object

Default: {}

lambda_dlq Example

```
"lambda_dlq": {  
    "TargetArn": "arn:aws:sns:us-east-1:accountnumber:topic"  
}
```

lambda_environment

Environment variables which are passed to the lambda function.

lambda_environment Keys

Variables : Dictionary of environment variables.

Type: object

Default: null

lambda_environment Example

```
{
    "lambda_environment": {
        "Variables": {
            "VAR1": "val1",
            "VAR2": "val2",
            "VAR3": "val3"
        }
    }
}
```

lambda_filesystems

List of Dictionaries that are passed with the EFS filesystem configuration. Expects the ARN of the filesystem and the Local Mount Path

Type: list
Default: []

lambda_filesystems Example

```
{
    "Arn": "arn",
    "LocalMountPath": "/mnt/efs/"
}
```

lambda_layers

List of AWS Lambda Layer ARNs to add to Lambda Function

Type: list
Default: []

lambda_memory

The amount of memory to give a Lambda function

Type: string
Default: "128"
Units: Megabytes

lambda_provisioned_throughput

This will allow provisioned throughput of a lambda function. This specifically will ensure the function is warmed for a provisioned amount to eliminate any function cold starts (not to be confused with VPC cold starts)

More info on provisioned throughput can be found here: <https://aws.amazon.com/blogs/aws/new-provisioned-concurrency-for-lambda-functions/>

Type: int

Default: null

lambda_role

Override the default generated IAM Role name.

Type: string

Default: "\${group}_\${app}_role"

lambda_subnet_count

Enables ability to specify subnet resiliency of lambda functions. By default, uses all subnets of type subnet_purpose specified.

Each VPC in your AWS account has a Hyperplane ENI limit. The ENI Limit determines how many Hyperplane ENIs you can have in one VPC. The limit applies to Lambda in the same VPC and is set to 250 by default. If you exceed a ENI Limit, Lambda deployment will fail with a Hyperplane ENI Limit error.

At this time, you will need to submit a limit increase or reduce how many SG:Subnet Tuples you have per function. When you connect a function to a VPC, Lambda creates an elastic network interface for each combination of security group and subnet in your function's VPC configuration.

More info on limits can be found here: <https://docs.aws.amazon.com/lambda/latest/dg/limits.html>

Type: int

Default: <>MAX SUBNET COUNT>>

lambda_subnet_purpose

Determines if the AWS Lambda should be public (external) or non-public (internal).

Type: string

Default: "internal"

Options

- "internal"
- "external"

lambda_timeout

The timeout setting for Lambda function. See official limits <https://docs.aws.amazon.com/lambda/latest/dg/limits.html>

Type: string

Default: "900"

Units: Seconds

lambda_tracing

Lambda Tracing feature allows you to enable X-Ray APIs to your lambda function to identify performance bottlenecks and troubleshoot requests that are in error.

If you've enabled X-Ray tracing in a service that invokes your function, Lambda sends traces to X-Ray automatically. The upstream service, such as Amazon API Gateway, or an application hosted on Amazon EC2 that is instrumented with the X-Ray SDK, samples incoming requests and adds a tracing header that tells Lambda to send traces or not. For a full list of services that support active instrumentation, see Supported AWS Services in the AWS X-Ray Developer Guide. For more details see: <https://docs.aws.amazon.com/lambda/latest/dg/lambda-x-ray.html>

Currently AWS API supports Active or PassThrough.

Type: Object

Default: {}

lambda_tracing Example

```
"lambda_tracing": {  
    "Mode": "Active"  
}
```

cloudfunction_timeout

The function execution timeout. Execution is considered failed and can be terminated if the function is not completed at the end of the timeout period. Defaults to 60 seconds.

A duration in seconds with up to nine fractional digits, terminated by ‘s’. Example: “3.5s”.

Type: String

Default: None

Example: "60s"

cloudfunction_memory_mb

Memory in Mb specified as an integer (without Mb or mb after the value). GCP currently defaults to 256Mb if no value is given.

Type: Integer

Default: None

Example: 128

cloudfunction_environment

Environment variables that should be present when the Cloud Function is invoked.

Type: Dictionary

Default: None

Example: { 'MY_ENV_VAR': 'My value!' }

cloudfunction_allow_unauthenticated

Creates an IAM Binding which will allow anonymous/unauthenticated users to invoke the function. This only applies to HTTP triggers, event triggers cannot be invoked by users directly. This option mimics the gcloud functions deploy ... --allow-unauthenticated CLI option and adds an IAM binding for ‘allUsers’ with ‘roles/cloudfunctions.invoker’.

Type: Boolean

Default: False

cloudfunction_iam_bindings

Updates the Cloud Function's IAM Policy, which can be used to control which users or service accounts can access the function. It allows granular control on who has permissions to the function, and not which permissions the function itself has. For allowing anonymous access to the function the `cloudfunction_allow_unauthenticated=True` option is simpler.

Type: Array

Default: []

Example:

```
[  
  {  
    "members": [  
      "user:jon.snow@GameOfThrones.com",  
      "serviceAccount:my-service-acccount@my-project.iam.  
      ↵gserviceaccount.com",  
      ],  
      "role": "roles/cloudfunctions.invoker"  
    }  
]
```

Note:

The role `roles/cloudfunctions.invoker` does not allow invoking via `gcloud functions call ...`, instead use a command like

```
curl $URL -H "Authorization: bearer $(gcloud auth print-identity-token)" to  
test granular invocation permissions
```

cloudfunction_max_instances

Maximum number of instances of a function that can run in parallel. GCP defaults to no limit if a value is not given.

Type: Integer

Default: None

Example: 5

cloudfunction_ingress_type

Ingress type to use. Foremast does not have a default, however GCP Defaults to ALLOW_ALL if none is given. Options are: INGRESS_SETTINGS_UNSPECIFIED, ALLOW_ALL, ALLOW_INTERNAL_ONLY, ALLOW_INTERNAL_AND_GCLB

For information on this option see the [GCP Documentation on Ingress Settings](#).

Type: String

Default: None

Example: ALLOW_INTERNAL_ONLY

cloudfunction_vpc

```
"cloudfunction_vpc": {
    "connector": {
        "us-central1": "projects/your-vpc-project/locations/us-central1/connectors/stage-
        ↵us-central1",
        "us-east1": "projects/your-vpc-project/locations/us-east1/connectors/stage-us-
        ↵east1"
    },
    "egress_type": "PRIVATE_RANGES_ONLY"
}
```

connector

VPC Connector to use, which will allow private VPC network access to the Cloud Function. Should be defined as key/value pairs where the key is the region and the value is the VPC connector.

Type: Dictionary

Default: None

Example:

```
{
    "us-central1": "projects/your-host-project/locations/us-central1/
    ↵connectors/yourconnector-us-central1",
    "us-east1": "projects/your-host-project/locations/us-east1/
    ↵connectors/yourconnector-us-east1"
}
```

egress_type

Egress type to use. Foremast does not have a default, however GCP Defaults to PRIVATE_RANGES_ONLY if none is given. Options are: VPC_CONNECTOR_EGRESS_SETTINGS_UNSPECIFIED, PRIVATE_RANGES_ONLY, ALL_TRAFFIC.

For information on this option see the [GCP Documentation on VPC Egress Settings](#).

Type: String

Default: None

Example: PRIVATE_RANGES_ONLY

cloudfunction_event_trigger

Configures a trigger for a GCP Cloud Function. If none is given, GCP will default to an HTTPS trigger. Trigger types are immutable in GCP, so once a trigger type is used (https, pub/sub, GCS, etc) it cannot be changed. It is possible to change the resource used in the trigger, but not the trigger type itself.

Example Pub/Sub trigger:

```
"cloudfunction_event_trigger": {
    "event_type": "providers/cloud.pubsub/eventTypes/topic.publish",
    "resource": "/topics/my_topic",
    "failure_policy": {
```

(continues on next page)

(continued from previous page)

```
        "retry": true
    }
}
```

Example GCS Bucket trigger:

```
"cloudfunction_event_trigger": {
    "resource": "buckets/my_bucket_name",
    "event_type": "google.storage.object.archive",
    "failure_policy": {
        "retry": false
    }
}
```

event_type

Event type to trigger the Cloud Function. Event types and their formats can vary, the easiest way to determine your event type is to run the command `gcloud functions event-types list`. and refer to the EVENT_TYPE column.

Type: String

Default: None

Example Pub/Sub: providers/cloud.pubsub/eventTypes/topic.publish

Example Storage: google.pubsub.topic.publish

Example Firestore Storage: providers/cloud.firestore/eventTypes/document.write

resource

The resource to trigger off of. The resource type given must match the event_type specified. For example, a resource path to a GCS Bucket with a Pub/Sub event trigger will be rejected. GCP expects the project to be specified and the full path to the resource, however if omitted Foremast will add this automatically.

Type: String

Default: None

Example Pub/Sub: topics/my_topic

Example Storage: buckets/my_bucket

asg Block

Top level key containing information regarding application ASGs

hc_type

Note: See `foremast.pipeline.construct_pipeline_block.construct_pipeline_block()` for cases where the Health Check type is overridden to "EC2".

ASG Health check type (EC2 or ELB)

Type: string

Default: "ELB"

Options:

- "ELB"
- "EC2"

app_grace_period

App specific health check grace period (added onto default ASG healthcheck grace period) to delay sending of health check requests. This is useful in the event your application takes longer to boot than the default hc_grace_period defined in templates.

For example, hc_grace_period may be 180 seconds, but an app may need a variable amount of time to boot (say 30 seconds extra). This will add 180 + 30 to calculate the overall hc_grace_period of 210 seconds.

Type: number

Default: 0

Units: Seconds

max_inst

Maximum number of instances ASG will scale to.

Type: number

Default: 3

min_inst

Minimum number of instances your auto-scaling group should have at all times. This is also the default number of instances

Type: number

Default: 1

ssh_keypair

SSH key that your EC2 instances will use. Must already be created in AWS. This replaces the non-functional and deprecated app_ssh_key configuration key.

Type: string

Default: "{{ account }}_{{ region }}_default" - {{ account }} being the AWS account in the configuration name

subnet_purpose

Determines if the instances should be public (external) or non-public (internal).

Type: string

Default: "internal"

Options

- "internal"
- "external"

enable_public_ips

Determines if instances in an cluster should have public IPs associated. By default, this is set to null which means it uses default behavior configured for your subnets in your cloud provider.

Type: boolean

Default: null

Options

- true
- false

scaling_policy

To better explain this feature, this has been moved to: [scaling_policy - V1 Cluster Scaling](#)

custom_scaling_policies

To better explain this feature, this has been moved to: [custom_scaling_policies - V2 Cluster Scaling](#)

scheduled_actions

To better explain this feature, this has been moved to: [scheduled_actions](#)

elb Block

Top level key for ELB configuration

access_log

Access Log configuration block. Ensure S3 bucket has proper bucket policy to enable writing.

access_log Keys

bucket_name : Name of S3 bucket to write access log to

Type: string

Default: Null

bucket_prefix : Prefix to write to in the S3 bucket

Type: string

Default: Null

`emit_interval`: ELB Access Log write delay

Type: number

Range: 5-60

Units: seconds

Default: Null

`connection_draining_timeout`

Connection Draining Timeout to set on the ELB. This allows existing requests to complete before the load balancer shifts traffic away from a deregistered or unhealthy instance.

Type: number

Range: 1-3600

Units: seconds

Default: Null

`certificate`

Name of SSL certification for ELB. SSL certificate must be uploaded to AWS first.

Type: string

Default: Null

`health`

Health check configuration block

`health Keys`

`interval`: ELB health check interval

Type: number

Units: seconds

Default: 20

`threshold`: Number of consecutive health check successes before declaring EC2 instance healthy.

Type: number

Default: 2

`timeout`: Health check response timeout

Type: number

Units: seconds

Default: 10

`unhealthy_threshold`: number of consecutive health check failures before declaring EC2 instance unhealthy

Type: number

Default: 5

idle_timeout

Idle Timeout to set on the ELB. This is the time, in seconds, that the connection is allowed to be idle (no data has been sent over the connection) before it is closed by the load balancer.

Type: number

Range: 1-3600

Units: seconds

Default: 60

ports

Defines ELB listeners. Expects a list of listeners.

ports Keys

`instance` : The protocol:port of the instance

Type: string

Default: "HTTP:8080"

`loadbalancer` : the protocol:port of the load balancer

Type: string

Default: "HTTP:80"

`stickiness` : defines stickiness on ELB; if app, specify cookie_name, if elb, specify cookie_ttl

Type: object

Default: None

Supported Types: elb, app

Example app:

```
{  
    "stickiness": {  
        "type": "app",  
        "cookie_name": "$cookiename"  
    }  
}
```

Example elb:

```
{  
    "stickiness": {  
        "type": "elb",  
        "cookie_ttl": 300  
    }  
}
```

`certificate` : The name of the certificate to use if required

Type: string

Default: null

`listener_policies` : A list of listener policies to associate to an ELB. Must be created in AWS first.

Type: array

Default: []

`backend_policies` : A list of backend server policies to associate to an ELB. Must be created in AWS first.

Type: array

Default: []

Example: ["WebSocket-Proxy-Protocol"]`

ports Example

```
{
    "ports": [
        {
            "instance": "HTTP:8080",
            "loadbalancer": "HTTP:80",
            "stickiness": {
                "type": "app",
                "cookie_name": "cookie"
            }
        },
        {
            "certificate": "my_cert",
            "instance": "HTTP:8443",
            "loadbalancer": "HTTPS:443",
            "listener_policies": [
                "MyExamplePolicy"
            ],
            "stickiness": {
                "type": "elb",
                "cookie_name": 300
            }
        }
    ]
}
```

subnet_purpose

Determines if the load balancer should be public (external) or non-public (internal). When changing this option, the ELB and DNS Records must be manually destroyed before deployment. This is necessary because the ELB Scheme is not modifiable.

Type: string

Default: "internal"

Options:

- "internal"
- "external"

target

The check the ELB will use to validate application is online.

Type: string

Default: "TCP :8080"

regions Key

Dictionary of AWS regions that application will be deployed to.

Type: array

Default: { "us-east-1": {} }

deploy_strategy Key

Spinnaker strategy to use for deployments.

Type: string

Default: "highlander"

Options:

- "highlander" - destroy old server group
- "redblack" - disables old server group but do not destroy
- "canary" - Only used in S3 deployments. Causes pipeline to first deploy to CANARY path
- "alpha" - Only used in S3 deployments. Causes pipeline to first deploy to an ALPHA path
- "mirror" - Only used in S3 deployments. Contents are deployed as-is, no version or LATEST directory
- "branchrelease" - Only used in S3 deployments. S3 Folders correlate to Git Branches, using versions and LATEST directory

security_group Block

Hold configuration for creating application specific security group

description

Description of the security group. Used in AWS for creation

Type: string

Default: "Auto-Gen SG for {{ app }}"

elb_extras

A list of extra security groups to assign to ELB

Type: array

Default: []

Example: ["all_access", "test_sg"]`

instance_extras

A list of extra security groups to assign to each instance

Type: array

Default: []

Example: ["all_access", "test_sg"]`

ingress

Provides a list of other security groups and ports to allow inbound access to application

egress

Provides info about outbound access from application

Type: string

Default: "0.0.0.0/0" `

security_group Example

You can reference SG by name or by cidr block, you can also specify cross account SG by name by referring to the spinnaker environment name. To see an example of this see below:

```
{
    "security_group": {
        "ingress": {
            "examplesecuritygroupname": [
                {"start_port": 80, "end_port": 80, "protocol": "tcp"},
                {"start_port": 443, "end_port": 443, "protocol": "tcp"},
                {"start_port": 443, "end_port": 443, "protocol": "tcp", "env": "prod"}
            ],
            "192.168.100.0/24": [
                {"start_port": 80, "end_port": 80, "protocol": "tcp"}
            ]
        },
        "egress": {
            "192.168.100.0/24": [
                {"start_port": 80, "end_port": 80, "protocol": "tcp"}
            ]
        }
    }
}
```

dns Block

Top level key for dns settings

ttl

Defines DNS TTL for generated DNS records

Type: number
Units: seconds
Default: 60

lambda_triggers

A list of all events to trigger a Lambda function. See [Lambda Triggers and Events](#) for details

Type: array
Default: []

datapipeline Block

Top level key for AWS Data Pipeline settings. Only necessary for Data Pipeline deployments.

name

Name of the Data Pipeline. This defaults to the application name.

Type: string
Default: \$appname

description

Description of the Data Pipeline.

Type: string
Default: ""

activate_on_deploy

Activates a Data Pipeline after deployment. Useful for OnDemand pipelines

Type: boolean
Default: false

json_definition

The exported JSON definition of the AWS Data Pipeline. You can get this by clicking “Export” in the AWS Console when creating the Data Pipeline.

Type: object
Default: {}

s3 Block

Holds settings related to s3 deployments

path

Path to upload assets to in a specified s3 bucket. Only works for S3 pipelines not using shared/master bucket setup. Refer to *s3_bucket_master* for more information.

Type: string

Default: "/"

bucket_acl

General ACL to apply to S3 bucket

Type: string

Default: "private"

Options:

- "public"
- "private"

bucket_name

Allows an S3 bucket name to be specified vs generated by pipeline

Type: string

Default: ""

bucket_policy

The S3 bucket policy in json format to apply to created S3 bucket. Must be a valid S3 bucket policy; use the AWS policy generator/simulator to test your policy. (<https://awspolicygen.s3.amazonaws.com/policygen.html>)

Type: json

Default: "{}"

content_metadata

S3 object metadata based on path. The "path" field should have NO leading or trailing slashes.

Type: object

Default: None

Example config:

```
[  
  {  
    "path": "assets/compressed",  
    "content-encoding": "br"  
  },  
  {  
    "path": "assets/gzip",  
    "content-encoding": "gzip"  
  }]
```

shared_bucket_master

Setups up an S3 bucket as a shared target so other Spinnaker pipelines can upload to it. i

Type: boolean

Default: false

Example:

An example of this is having one s3 bucket for a given website. This website has a commercial and business webpage that are two unique deployment pipelines. Each of the unique apps would specify the *shared_bucket_target* to the Spinnaker application name of the *shared_bucket_master* pipeline. It is common to have a bare source repository for a master bucket with just Foremast pipeline and application configuration files.

shared_bucket_target

Shared bucket to deploy to. Refer to *shared_bucket_master* for use case and example.

Type: string

Default: None

cors

S3 CORS configuration block

cors Keys

`enabled`: Enables/Disables CORS configuration

Type: boolean

Default: false

`cors_rules`:

A list of CORS rules including lists of headers, methods, origins, exposed headers, and max age. For more details refer to: http://boto3.readthedocs.io/en/latest/reference/services/s3.html#S3.Client.put_bucket_cors

Type: object

Default: None

Example config:

```
[  
  {  
    "cors_headers": [],  
    "cors_methods": [],  
    "cors_origins": [],  
    "cors_expose_headers": [],  
    "cors_max_age": 600  
  },  
  {  
    "cors_headers": [],  
    "cors_methods": [],  
    "cors_origins": [],  
    "cors_expose_headers": []  
  }]
```

(continues on next page)

(continued from previous page)

```

        "cors_max_age": 600
    }
]
```

encryption

S3 Encryption configuration block

encryption Keys

`enabled` : Enables/Disables S3 Encryption configuration

Type: boolean

Default: false

`encryption_rules`:

A list of S3 encryption rules. As of today only one rule is supported: `ApplyServerSideEncryptionByDefault`. Built in support for additional rules if this changes. Support for both AES256 or custom KMS (`aws:kms`) `SSEAlgorithm`. For `aws:kms`, specify a custom `KMSMasterKeyID`; this is not needed for AES256 and should not be specified. For more details refer to: http://boto3.readthedocs.io/en/latest/reference/services/s3.html#S3.Client.put_bucket_encryption

Type: object

Default: None

Example config:

```

[
    {
        "ApplyServerSideEncryptionByDefault": {
            'SSEAlgorithm': 'AES256'|'aws:kms',
            'KMSMasterKeyID': 'string'
        }
    }
]
```

lifecycle

S3 Lifecycle configuration block

lifecycle Keys

`enabled` : Enables/Disables S3 Lifecycle configuration

Type: boolean

Default: false

`lifecycle_rules`:

A list of S3 lifecycle rules, if a lifecycle exists it replaces it. For more details refer to: http://boto3.readthedocs.io/en/latest/reference/services/s3.html#S3.Client.put_bucket_lifecycle_configuration

Type: array

Default: [{}]

Example config:

```
[  
  {  
    'Expiration': {  
      'Date': datetime(2015, 1, 1),  
      'Days': 123,  
      'ExpiredObjectDeleteMarker': True|False  
    },  
    'ID': 'string',  
    'Prefix': 'string',  
    'Filter': {  
      'Prefix': 'string',  
      'Tag': {  
        'Key': 'string',  
        'Value': 'string'  
      },  
      'And': {  
        'Prefix': 'string',  
        'Tags': [  
          {  
            'Key': 'string',  
            'Value': 'string'  
          },  
          ...  
        ]  
      }  
    },  
    'Status': 'Enabled'|'Disabled',  
    'Transitions': [  
      {  
        'Date': datetime(2015, 1, 1),  
        'Days': 123,  
        'StorageClass': 'GLACIER'|'STANDARD_IA'|'ONEZONE_IA'  
      },  
      ...  
    ],  
    'NoncurrentVersionTransitions': [  
      {  
        'NoncurrentDays': 123,  
        'StorageClass': 'GLACIER'|'STANDARD_IA'|'ONEZONE_IA'  
      },  
      ...  
    ],  
    'NoncurrentVersionExpiration': {  
      'NoncurrentDays': 123  
    },  
    'AbortIncompleteMultipartUpload': {  
      'DaysAfterInitiation': 123  
    }  
  }  
]
```

logging

S3 access logging configuration block

logging Keys

`enabled`: Enables/Disables S3 logging configuration

Type: boolean

Default: false

`logging_grants`:

Specify permissions for who can view and modify the logging parameters. To set the logging status of a bucket, you must be the bucket owner. For more details refer to: http://boto3.readthedocs.io/en/latest/reference/services/s3.html#S3.Client.put_bucket_logging

Type: object

Default: None

Example config:

```
[  
  {  
    'Grantee': {  
      'DisplayName': 'string',  
      'EmailAddress': 'string',  
      'ID': 'string',  
      'Type': 'CanonicalUser'|'AmazonCustomerByEmail'|'Group',  
      'URI': 'string'  
    },  
    'Permission': 'FULL_CONTROL'|'READ'|'WRITE'  
  }  
]
```

`logging_bucket` : Specifies the bucket where you want Amazon S3 to store server access logs.

Type: string

Default: ""

`logging_bucket_prefix` : This element lets you specify a prefix for the keys that the log files will be stored under.

Type: string

Default: {{ app }} /

notification

S3 Notification configuration block

notification Keys

`enabled`: Enables/Disables S3 Notification configuration

Type: boolean

Default: false

`topic_configurations`:

A list of S3 SNS topic notification rules, if an SNS notification configuration exists it replaces it. For more details refer to: https://boto3.amazonaws.com/v1/documentation/api/latest/reference/services/s3.html#S3.Client.put_bucket_notification_configuration

Type: array

Default: [{}]

Example config:

```
queue_configurations:
```

A list of S3 SQS notification rules, if an SQS notification configuration exists it replaces it. For more details refer to: https://boto3.amazonaws.com/v1/documentation/api/latest/reference/services/s3.html#S3.Client.put_bucket_notification_configuration

Type: array

Default: [{}]

Example config:

```
lambda_configurations:
```

A list of S3 Lambda notification rules, if an Lambda notification configuration exists it replaces it. For more details refer to: https://boto3.amazonaws.com/v1/documentation/api/latest/reference/services/s3.html#S3.Client.put_bucket_notification_configuration

Type: array

Default: [{}]

Example config:

tagging

S3 tagging configuration block

tagging Keys

`tags` : This element lets you specify a prefix for the keys that the log files will be stored under.

Type: object

Default: { }

Example config:

```
{  
    "key": "value",  
    "key1": "value1"  
}
```

versioning

S3 versioning configuration block

versioning Keys

enabled : Enables/Disables S3 versioning configuration

Type: boolean

Default: false

mfa_delete : Specifies whether MFA delete is enabled in the bucket versioning configuration.

Type: string

Default: Disabled

website

S3 Website configuration block

website Keys

enabled : Enables/Disables an S3 bucket from being website enabled

Type: boolean

Default: true

index_suffix : Default index page

Type: string

Default: "index.html"

error_document : Default error page

Type: string

Default: "404.html"

stepfunction Block

Top level key for AWS Step Function settings. Only necessary for Step Function deployments.

statemachine_type

Determines whether a Standard or Express state machine is created. You cannot update the type of a state machine once it has been created.

Type: string

Default: "STANDARD"

Options:

- "STANDARD"
- "EXPRESS"

json_definition

The exported JSON definition of the AWS Step Function State Machine. You could craft this using the GUI or by using the AWS Toolkit in various IDEs.

Type: object

Default: { }

tracing

Determine whether AWS X-Ray tracing is enabled.

Type: object

Default: { "enabled": false}

Options:

- { "enabled": false}
- { "enabled": true}

logging_configuration

Description of the Data Pipeline.

Type: object

Default: { }

Example:

```
{
    "level": "ALL"|"ERROR"|"FATAL"|"OFF",
    "includeExecutionData": true|false,
    "destinations": [
        {
            "cloudWatchLogsLogGroup": {
                "logGroupArn": "arn:aws:logs:region:account-id:log-
                ↴group:log_group_name"
            }
        }
    ]
}
```

qe Block

Top level key for Quality Test settings; used by Quality Testing Stages typically post and pre deployment. This specific feature is left up for custom Jinja Templates as most have specific testing flows. All keys in the QE block are passed and can be interpreted by custom Jinja2 Stages. While we are providing a base example below, every user of foremast might need to tweak the logic to add/remove what they need in their respective organization!

To facilitate this concept, this section highlights some common keys and ideas that can be implemented to achieve testing Post Deploy stages (using custom parameters per environment).

Refer to our full example templates here: <https://github.com/foremast/foremast-template-examples>

failure_action

Define how Spinnaker should handle a quality stage failure. #. *fail_pipeline* - default behaviour #. *fail_branch* - only fail that branch of the pipeline #. *fail_branch_continue_pipeline* - fail the branch but continue the pipeline; then fail the pipeline #. *ignore_failures* - ignores issues

Type: string

Default: "fail_pipeline"

Values:

- "fail_pipeline"
- "fail_branch"
- "fail_branch_continue_pipeline"
- "ignore_failures"

ignore_unstable_results

If set to true, ignore and treat unstable results from Jenkins stage as success.

Type: boolean

jenkins_master

Name of Jenkins Master configured in Spinnaker to run your Jobs against

Type: string

test_job

Name of Jenkins Job to Run. This defaults to the application name.

Type: string

test_params

Dictionary of Jenkins Job Parameters K/V pairs to pass down to jobs.

Type: dict

stage_timeout

Timeout quality tests step after a specified amount of time in milliseconds

Type: int

Format: ms

wait_for_completion

If set to true, wait until the Jenkins stage is complete to proceed.

Type: boolean

Example Quality Stage JSON Configuration

```
{  
    "requisiteStageRefIds": [""],  
    "refId": "master",  
    "type": "jenkins",  
    "name": "{{ data.app.environment|upper }}: Quality Tests",  
    "waitForCompletion": {{ data.qe.wait_for_completion|default(true)|tojson }},  
    "markUnstableAsSuccessful": {{ data.qe.ignore_unstable_  
→results|default(false)|tojson }},  
    "master": "{{ data.qe.jenkins_master or 'jenkinsci' }}",  
    {%- if data.qe.stage_timeout %}  
        "stageTimeoutMs": {{ data.qe.stage_timeout }},  
    {%- endif %}  
    {%- if data.qe %}  
        "job": "{{ data.qe.test_job }}",  
        {%- if data.qe.stage_failure_action %}  
            {%- if data.qe.stage_failure_action == "fail_pipeline" %}  
                "completeOtherBranchesThenFail": false,  
                "continuePipeline": false,  
                "failPipeline": true,  
            {%- elif data.qe.stage_failure_action == "fail_branch" %}  
                "completeOtherBranchesThenFail": false,  
                "continuePipeline": false,  
                "failPipeline": false,  
            {%- elif data.qe.stage_failure_action == "fail_branch_continue_pipeline" %}  
                "completeOtherBranchesThenFail": true,  
                "continuePipeline": false,  
                "failPipeline": false,  
            {%- elif data.qe.stage_failure_action == "ignore_failures" %}  
                "completeOtherBranchesThenFail": false,  
                "continuePipeline": true,  
                "failPipeline": false,  
            {%- endif %}  
        {%- else %}  
            "completeOtherBranchesThenFail": false,  
            "continuePipeline": false,  
            "failPipeline": true,  
        {%- endif %}  
        "parameters": {  
            {%- if data.qe.test_params %}  
                {%- for param_name, param_value in data.qe.test_params.items() %}  
                    "{{ param_name }}": "{{ param_value }}"{%- if not loop.last -%},{%-  
→endif -%}  
                {%- endfor %}  
            {%- endif %}  
            {%- if data.qe.test_json %}  
                {%- if data.qe.test_params %},{%- endif %}  
                "test_json": "{{ data.qe }}"  
            {%- endif %}  
        }  
    }  
}
```

(continues on next page)

(continued from previous page)

```

        }
    {%
      else %
        "completeOtherBranchesThenFail": false,
        "continuePipeline": false,
        "failPipeline": true,
        "job": "spinnaker-qe-{{ data.app.environment }}",
        "parameters": {
          "QE_LEVEL": "{{ data.qe.test_type or 'load' }}",
          "SPINNAKER_APP_NAME": "{{ data.app.appname }}"
        }
    {%
      endif %
  }
}

```

Completion Webhooks are Spinnaker Webhook Stages that are appended to the pipeline stages for this environment.

```
{
  "completion_webhooks": [
    {
      "url": "https://webhook.com/webhook1",
      "custom_headers": {
        "my-header": "hello"
      },
      "method": "POST",
      "name": "Webhook 1",
      "payload": {
        "webhook": "one"
      }
    },
    {
      "url": "https://webhook.com/webhook2",
      "custom_headers": {
        "my-header": "hello again"
      },
      "method": "POST",
      "name": "Webhook 2",
      "payload": {
        "webhook": "two"
      }
    }
  ]
}
```

1.3.4 Region Specific Overrides

- *Purpose*
- *Example*

Purpose

Within the *application.json* configuration, the need may arise to use different settings for different regions. You can override any setting in the *regions* blocks and that will be applied to only a specific region.

Example

```
{  
    "security_group": {  
        "description": "something useful",  
        "elb_extras": [],  
        "instance_extras": ["offices_all"]  
    },  
    "app": {  
        "instance_type": "t2.small",  
        "app_description": "Edge Forrest Demo application",  
        "instance_profile": "forrest_edge_profile"  
    },  
    "elb": {  
        "subnet_purpose": "internal",  
        "target": "TCP:8080",  
        "ports": [  
            {"loadbalancer": "HTTP:80", "instance": "HTTP:8080"}  
        ]  
    },  
    "asg": {  
        "subnet_purpose": "internal",  
        "min_inst": 1,  
        "max_inst": 1  
    },  
    "dns": { "ttl": 120 },  
    "regions": {  
        "us-east-1": {},  
        "us-west-2": {  
            "app": {  
                "instance_type": "t2.medium"  
            },  
            "asg": {  
                "min_inst": 5,  
                "max_inst": 10  
            }  
        }  
    }  
}
```

In the above example, under the `regions` blocks region-specific configs are set for `us-west-2`. These configs override what is in the main json block. `us-east-1` just has an empty `{}` and so no settings are specifically overridden and it will just use values from the main json block.

The empty `{}` is necessary for any regions without overrides. If you did not include `"us-east-1": {}` in the above example, the application would only deploy the `us-west-2`.

1.3.5 ami-lookup.json

- *Purpose*
- *Example Json*
- *Json Location*

Purpose

This json file is used as an AMI ID look up table for each region. It is used during the bake stage of Spinnaker deployments in order to determine the base AMI ID to use for baking.

Example Json

```
{
    "us-east-1": {
        "origin": "ami-xxxx",
        "origin_default": "ami-xxxx",
        "origin_fedora": "ami-xxxx",
        "origin_amazon": "ami-xxxx",
        "origin_ubuntu": "ami-xxxx",
        "origin_debian": "ami-xxxx",
        "origin_testing": "ami-xxxx",
    }
    "us-west-2": {
        "origin": "ami-xxxx",
        "origin_default": "ami-xxxx",
        "origin_fedora": "ami-xxxx",
        "origin_amazon": "ami-xxxx",
        "origin_ubuntu": "ami-xxxx",
        "origin_debian": "ami-xxxx",
        "origin_testing": "ami-xxxx",
    }
}
```

Json Location

Foremast will look for this information at `ami_json_url` defined in `foremast.cfg / config.py`. For example, you can host the file named `ami-lookup.json` in an S3 bucket and then set `ami_json_url = http://s3bucketurl.com/ami-lookup.json`.

You can host this file anywhere as long as an HTTP GET will return the JSON and a 2XX.

1.3.6 AWS Credentials

- *Purpose*
- *Example Configuration*
- *Configuration Location*
- *Configuration Details*

Purpose

This is how AWS credentials are stored for usage with Foremast. All AWS calls outside of Spinnaker use Boto3 so standard Boto3 locations work but account/environment must be specified.

Example Configuration

```
[build]
aws_access_key_id = XXXXXXXXXXXXXXXXXXXXXXXXX
aws_secret_access_key = yyyyyyyyyxxxxxxxxyyyyyyyyyyyyxxxxxx

[dev]
aws_access_key_id = AAAAAAAAAAAAAAAA
aws_secret_access_key = bbbbbbbbaaaaaaaaaaaaaabbbbbbbbbbbaaaaa

[stage]
aws_access_key_id = TTTTTTTTTTTTTTTTTTTTTTTTTTT
aws_secret_access_key = ssssssssssssttttttttttssssssss
```

Configuration Location

Foremast just uses Boto3 which will look at `~/.aws/credentials` for the credentials file.

Configuration Details

This is a standard Boto3 credentials file. You can read more about it on the [Boto3 docs](#). The important part is that each account/environment that Foremast is managing has a distinct section in `credentials`.

1.3.7 GCP Credentials

- *Purpose*
- *Example Configuration*
- *Service Account Basics*
- *Service Account Permissions*

Purpose

This is how AWS credentials are stored for usage with Foremast. All AWS calls outside of Spinnaker use Boto3 so standard Boto3 locations work but account/environment must be specified.

This section explains how GCP credentials are stored for usage with Foremast. All GCP calls outside of Spinnaker use the official GCP Python Clients and service account authentication.

Example Configuration

See the GCP Configuration section for details: [\[gcp\]](#)

Service Account Basics

Each GCP environment defined in Foremast needs a path to a json service account. You can export these from GCP's IAM Console. Foremast will only modify one GCP environment at a time, meaning when a pipeline is running for

env1, env2 will not be modified. **We recommend using a different service account per environment (with limited permissions) as an additional step to ensure security and proper segmentation of environments.**

Service Account Permissions

Permissions needed by Foremast can vary on a project-by-project basis. You can set these up for each individual project, or use folders and org level IAM policies to assign permissions across multiple projects.

Minimum Permissions:

Project IAM Admin (*roles/resourcemanager.projectIamAdmin*) is required in all projects for Foremast to get and set IAM permissions.

1.3.8 GCP Environments

- *Purpose*
- *Definition of GCP Environment*
- *GCP Project Labels*
 - *foremast_enabled*
 - *cloud_env*
 - *foremast_groups*
- *How to Label a Project in GCP*
- *Controlling Permissions on Projects*

Purpose

This document outlines how to define GCP environments in Foremast and the necessary configuration in GCP.

Definition of GCP Environment

Foremast uses environments with names dev, stage, prod, etc. when deploying and configuring applications. In AWS an environment maps to a single AWS account (or sub-account). GCP does not have a similar concept, and instead resources are organized within individual projects. Permissions can be managed at the project level (or higher using inheritance).

How projects in GCP are managed vary by organization. You may have a single project for each environment like yourcompany-prod and yourcompany-stage, or you may define projects by teams like yourcompany-yourteam-prod yourcompany-yourteam-stage. Further, you may have a strict project naming convention or no convention at all. You may use folders to define environments, teams, or no folders at all.

To support any combination of these variables Foremast defines an environment in GCP as a grouping of one or more projects using GCP project labels. GCP allows key/value pairs to be added to any project as a label, which Foremast can then easily query.

GCP Project Labels

`foremast_enabled`

When true, Foremast will consider this project when updating permissions and deploying applications.

Default: None

Required: Yes

Values: true or false

`cloud_env`

The name of the environment this project belongs to.

Default: None

Required: Yes

Values: Any string value, must match a GCP Environment defined in Foremast

`foremast_groups`

A __ (double underscore) separated list of groups that can request permissions to this project. (See [Controlling Permissions on Projects](#) for details)

Default: None

Required: No

Values: Double underscore separated list of Gitlab or GitHub groups

How to Label a Project in GCP

Labeling a project is done via the GCP Console or `gcloud` CLI tool. See the [GCP docs](#) on managing projects for details.

Controlling Permissions on Projects

Foremast will not grant permissions between projects with different `cloud_env` values, but Foremast can grant permissions between projects within the same environment (and does by default).

To ensure certain apps deployed via Foremast cannot request permissions in certain projects, you can use the `foremast_groups` label and pass in a double underscore separated list of groups permitted to access the project.

For example, assume you have two groups in Gitlab or Github: purchasing and customersupport. Each team also has their own GCP Projects: purchasing-prod and customersupport-prod. If there is no valid use-case for customer-support applications to request permissions to the purchasing-prod project, the label `foremast_groups=purchasing` can be added. This ensures only applications in the group purchasing can request permissions on this GCP project. To support multiple groups, simply add a double underscore and the additional group names: `foremast_groups=purchasing__anothergroup__anothergroup2`. If a customersupport application requests permissions to purchasing-prod, Foremast will raise an exception before any permissions/IAM modifications are made.

This allows Foremast administrators to lock down IAM permissions during deployments when needed. If `foremast_groups` is not set or has an empty value, any deployment in Foremast can request permissions to the given project.

1.3.9 Configuration Files Advanced Usages

Manual Pipelines

While Foremast has great support for many Spinnaker deployment features, it is not without flaws. Most notably, Foremast struggles in a few areas:

1. Limited support outside of AWS based pipelines
2. Keeping up with new features released in Spinnaker
3. Pipelines, Deployment Flows, and Structure can be seen as opinionated at times

While the Foremast's templating engine built around [Jinja2](#) is rather extensible, there is a bit of boilerplate code that needs to be written to support custom pipelines. This leads to many Foremast users not being able to Foremast to support new and/or more complex requirements not defined within Foremast just yet.

Regardless, we are still left with a need for a solution to manually creating pipelines via the Spinnaker UI. As a result, we have support in Foremast to allow users to specify “*manual*” pipeline type.

Manual pipelines allow users to store Spinnaker Pipeline JSON in a *RUNWAY_DIR* and allow Foremast to create/manage Spinnaker Pipelines using native Spinnaker Pipeline JSON. In addition, we enable the ability to store the JSON body as a Jinja2 Template (*json.j2*), allowing users to pass custom variables defined in Foremast configuration files to override common fields in Spinnaker Pipeline JSON.

While not ideal, this helps create support for things otherwise not currently supported in Foremast such as Kubernetes, AWS ECS, Google Cloud Functions, etc. More importantly it helps solve some of the issues noted above:

1. Spinnaker JSON + Foremast Templating = PROFIT
2. Create via Spinnaker UI, store the raw Spinnaker Pipeline JSON, Foremast does the heavy lifting of management
3. Unopinionated and Foremast only manages creation of pipelines (only acts as a template engine if specified)

To enable manual pipelines, a few top level “*pipeline*” keys are needed.

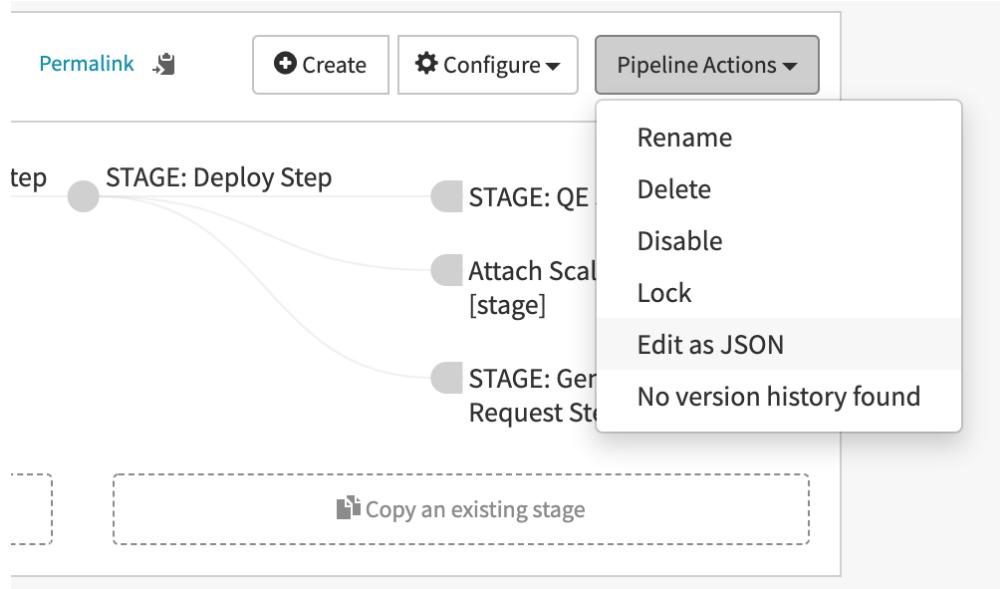
Manual Pipelines

Manual Pipelines allow users of Foremast to take raw Spinnaker Pipeline JSON and apply it via Foremast (thus version controlling their pipelines). This is useful for pipeline types in which Foremast is still adding support for. This is also an unopinionated interface to create templated Spinnaker pipelines using standard Jinja2.

Getting a Spinnaker Pipeline’s JSON Body

Spinnaker Pipelines are extremely customizable, extensible, and complex. The easiest way to create a Spinnaker Pipeline is to create one via the UI manually first!

1. Create your ideal pipeline via Spinnaker UI; valid, required stages, fully configured.
2. From the Pipeline Configuration Page, Click *Pipeline Actions* and select *Edit as JSON*.



3. Store JSON in a file in `RUNWAY_DIR`; name it whatever you need

Edit Pipeline JSON

The JSON below represents the pipeline configuration in its persisted state.

Note: Clicking "Update Pipeline" below will not save your changes to the server - it only updates the configuration within the browser, so you'll want to verify your changes and click "Save Changes" when you're ready.

```
1 * {
2   "executionEngine": "v3",
3   "keepWaitingPipelines": false,
4   "lastModifiedBy": "foremast",
5   "limitConcurrent": false,
6   "notifications": [],
7   "parallel": true,
8   "parameterConfig": [],
9   "stages": [
10     {
11       "comments": "application: job",
12       "failPipeline": true,
13       "instructions": "I confirm that I'm ready for this application to be promoted to <strong>stage</strong>.",
14       "judgmentInputs": [],
15       "name": "STAGE: Checkpoint",
16       "notifications": [],
17       "refId": "2",
18       "requisiteStageRefIds": [
19         "1"
20       ],
21       "type": "manualJudgment"
22     },
23   ],
24   "triggers": [
25     {
26       "enabled": true,
27       "job": "a_job",
28       "master": "jenkins",
29       "propertyFile": "vars",
30       "type": "jenkins"
31     }
32   ],
33   "updateTs": "1567626809000"
34 }
```

`pipeline.json Example`

Note:

The below example uses templated pipeline files managed inside of `TEMPLATES_PATH`. To specify a remote path, prefix a file name with: "`templates://`".

This enables Foremast users to specify common base templates for manual pipelines reducing duplicate templates in repositories.

In addition, this example is passing in some variables for use in the Jinja2 template using "template_variables".

```
{
    "deployment": "spinnaker",
    "type": "manual",
    "pipeline_files": [
        "templates://manual-pipeline.json.j2"
    ],
    "template_variables": [
        {
            "foo": "bar",
            "REGION": "us-west3"
        }
    ]
}
```

`pipeline.json` Keys

`pipeline_files`

Note:

Files could be of type `*.json` to denote a standard JSON document or files could be of type `*.json.j2` to denote Jinja2 template.

In addition, remote templates can be leverage by stating "`templates://`" URI. Foremast will look within the "`TEMPLATES_PATH`" specified in `config.py` for remote pipeline files.

List of JSON files to use for "manual" pipelines.

Type: Array

Default: []

`template_variables`

List of key/value pair objects to feed into templates. See below example on usage.

Type: Array of Objects

Default: []

Example Options:

`pipeline.json`

```
{
    "template_variables": [
        {
            "key": "value",
        }
    ]
}
```

(continues on next page)

(continued from previous page)

```
        "region": "us-west-2",
        "owner_email": "foo@example.com"
    }
]
}
```

manual-jinja-example.json.j2

```
[ {
    "schema" : "v2",
    "locked": {
        "allowUnlockUi": false,
        "ui": true
    },
    "protect": false,
    "metadata": {
        "name": "{{ template_variables.key }}",
        "description": "Deploys code to {{ template_",
        "owner": "{{ template_variables.name }}",
        "scopes": ["global"]
    }
    "pipeline": {},
    "triggers": []
}
]
```

Formatting your pipeline template file

Your pipeline template file should return an array with 1 or more Spinnaker pipeline definitions. Using an array allows Foremast to support the creation of more than one pipeline using a single pipeline template file. The most common usecase for this is creating two pipelines that are dependent on eachother, for example one pipeline that triggers when another finishes running. If only one pipeline is desired you should still use an array, but only place one Spinnaker pipeline definition in the array.

multiple-pipelines-jinja-example.json.j2

```
[ {
    "name" : "The first pipeline",
    "schema" : "v2",
    "locked": {
        "allowUnlockUi": false,
        "ui": true
    },
    "protect": false,
    "metadata": {
        "name": "{{ template_variables.key }}",
        "description": "Deploys code to {{ template_variables.region_",
        "owner": "{{ template_variables.name }}",
        "scopes": ["global"]
    }
    "pipeline": {},
    "triggers": []
},
{
    "name" : "The second pipeline",
}
```

(continues on next page)

(continued from previous page)

```

"schema" : "v2",
"locked": {
    "allowUnlockUi": false,
    "ui": true
},
"protect": false,
"metadata": {
    "name": "{{ template_variables.key }}",
    "description": "Deploys code to {{ template_variables.region }}",
    "owner": "{{ template_variables.name }}",
    "scopes": ["global"]
}
"pipeline": {},
"triggers": []
}
]

```

Note:

template_variables are shared per file. Multiple Spinnaker pipelines defined in a single file are sharing a common set of variables

Cluster Scaling Policies

Foremast Scaling Policies have two implementations:

1. A Foremast Managed Implementation (v1) - `scaling_policy`
2. A Custom Spinnaker Implementation (v2) - `custom_scaling_policies`

The intentions behind the scaling policy v1 implementation is that users of Foremast could simplify the complexities of scaling policies to end users. As a result, users could only specify their `scale_up` and `scale_down` fields.

With many recent enhancements and features in Spinnaker's API, we realized that some advanced users would prefer a more advanced implementation enabling teams to specify things like custom scaling policies leveraging custom metrics in their respective cloud provider. In addition, some teams have complex scaling policies that involve multiple steps. Finally, we wanted an unopinionated implementation that is true to the Spinnaker experience. As a result, the scaling policy v2 implementation was created.

In order to maintain support for both simple scaling policies, as well as advanced custom scaling policies, we have broken the implementation into two top level keys.

Note: When leveraging scaling policies in Foremast **only** one of the two implementations can be used.

If both `scaling_policy` and `custom_scaling_policies`, behavior will default to v1 `scaling_policy` for backwards compatibility.

`scaling_policy` - V1 Cluster Scaling

Defines scaling policy to attach to ASG. If this block does not exist, no scaling policy will be attached.

- *scaling_policy* Examples
 - Simple CPUUtilization Scaling Example
 - Custom Scale Up/Down Increments Example
- *scaling_policy* Keys
 - *metric*
 - *threshold*
 - *scale_down*
 - *increase_scaling_adjustment*
 - *decrease_scaling_adjustment*
 - *period_minutes*
 - *statistic*
 - *instance_warmup*

scaling_policy Examples

This section contains example usage but you are encouraged to modify and build your own scaling configurations that meet your needs.

Simple CPUUtilization Scaling Example

```
{  
    "scaling_policy": {  
        "metric": "CPUUtilization",  
        "threshold": 90,  
        "period_minutes": 10,  
        "instance_warmup": 180,  
        "statistic": "Average",  
        "scale_down": true  
    }  
}
```

Custom Scale Up/Down Increments Example

Note: To reduce nodes in a cluster the Scale Down API requires a negative number.

```
{  
    "scaling_policy": {  
        "metric": "CPUUtilization",  
        "threshold": 50,  
        "period_minutes": 1,  
        "instance_warmup": 120,  
        "scale_up": 1,  
        "scale_down": -1  
    }  
}
```

(continues on next page)

(continued from previous page)

```

    "statistic": "Average",
    "scale_down": true,
    "increase_scaling_adjustment": 5,
    "decrease_scaling_adjustment": -1
  }
}

```

scaling_policy Keys**metric**

The CloudWatch metric to trigger auto-scaling events.

Type: string

Default: "CPUUtilization"

Options:

- "CPUUtilization"
- "NetworkIn"
- "NetworkOut"
- "DiskReadBytes"

threshold

Metrics value limit for scaling up

Type: int

scale_down

Attach a default scale-down policy

Type: boolean

Default: true

increase_scaling_adjustment

Amount to increment by on scale up policies

Type: int

Default: 1

decrease_scaling_adjustment

Amount to decrement by on scale down policies. Negative numbers represent removing nodes from cluster.

Type: int

Default: -1

period_minutes

Time period to look across for determining if threshold was met. If you wish to have seconds, using a floating point such as .5 for 30 seconds.

Type: float

Default: 30

Units: Minutes

statistic

Statistic to calculate at the period to determine if threshold was met

Type: string

Default: "Average"

Options:

- "Average"
- "Maximum"
- "Minimum"
- "Sum"

instance_warmup

Time period to wait before adding metrics to Auto Scaling group

Type: int

Default: 600

Units: seconds

custom_scaling_policies - V2 Cluster Scaling

Warning: This is for advanced usage and expects understanding of how Spinnaker's API works.

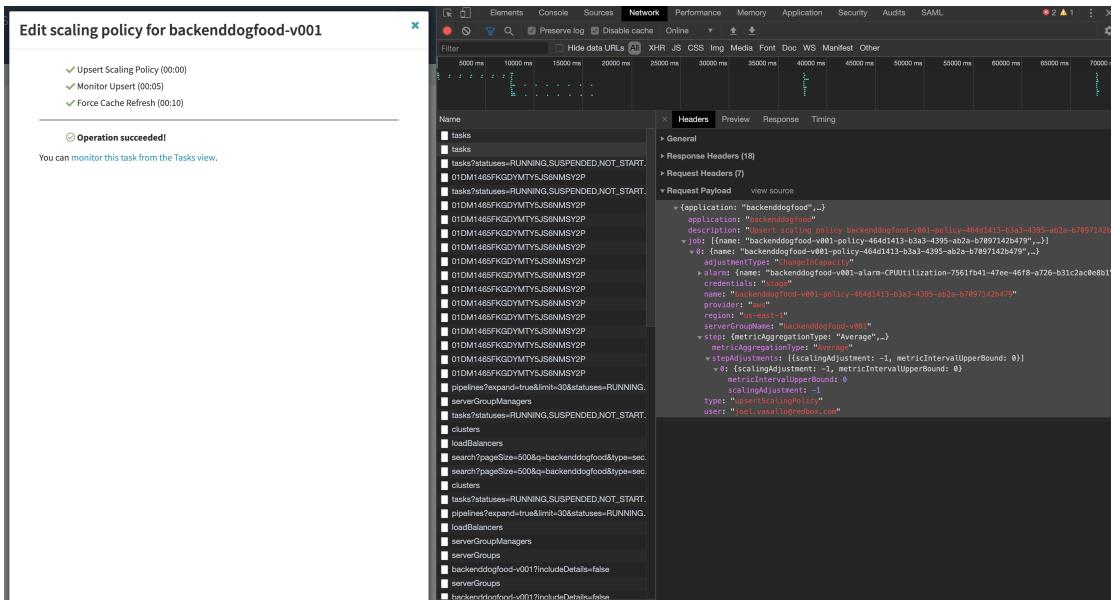
Enables the ability to define custom Spinnaker Cluster Scaling Policies, as defined by the Spinnaker API. This enables support for multiple scaling policies as well as custom metrics using Provider metrics. Currently, only tested with AWS AutoScaling groups.

There are no defaults set, however, we do have a helper method to lookup the latest AutoScalingGroupName if the dimension metric value is set to "\$self". This is mainly due to the fact the cluster name in Spinnaker is not known until deployment. In addition, all the templates are completely configurable if needed. Refer to the documentation section on using [Custom Pipelines](#) on using a custom templates_path.

- **Building/Creating `custom_scaling_policies`**
- `custom_scaling_policies` Examples
 - `custom_scaling_policies` Simple Step Scaling Example
 - `custom_scaling_policies` Multiple Scaling Policies Example
 - `custom_scaling_policies` Target Tracking Predefined Metrics Example
 - `custom_scaling_policies` Target Tracking Custom Metrics Example
- `custom_scaling_policies` Keys
 - `scaling_type`
 - `instance_warmup`
 - `disable_scale_in`
 - `scaling_metric` Keys
 - * `metric_type`
 - * `namespace`
 - * `metric_name`
 - * `statistic`
 - * `evaluation_periods`
 - * `evaluation_period`
 - * `threshold`
 - * `comparison_operator`
 - * `dimensions`
 - * `steps`

Building/Creating `custom_scaling_policies`

Custom values found below in our short examples and keys can be rather complex. For this reason, we recommend creating the scaling policies needed via the Spinnaker UI first and watching a browser network tab for the POST request body to the Spinnaker Gate "/tasks" API. This will contain the payload needed to get the values for the below usage example.



`custom_scaling_policies Examples`

This section contains example usage but you are encouraged to modify and build your own scaling configurations that meet your needs. Keys such as "namespace" and "metric_name" must be metrics from your provider and Spinnaker must have access to those metrics as well. The most common is AWS Cloudwatch Metrics. These values are hyper-specific to each end-user so not all keys will be defined in this document.

Note: To facilitate finding the name of cluster names generated by Spinnaker, users can specify a dimensions referencing "\$self" to let Foremast find the custom name

```
{
    "dimensions": [
        {
            "name": "AutoScalingGroupName",
            "value": "$self"
        }
    ],
    "custom_scaling_policies": [
        {
            "scaling_type": "step_scaling",
            "step_adjustments": [
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 75,
                    "upper_bound": 85,
                    "adjustment": 1
                },
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 85,
                    "upper_bound": 95,
                    "adjustment": 3
                },
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 95,
                    "upper_bound": 100,
                    "adjustment": 5
                }
            ],
            "evaluation_periods": 3,
            "metric_aggregation_type": "Average"
        }
    ]
}
```

`custom_scaling_policies Simple Step Scaling Example`

Note: This policy adds 1 instance between 75-85% CPU, 3 instances between 85-95, and 5 instances over 95% after 3 evaluation_periods of 1 minute.

```
{
    "custom_scaling_policies": [
        {
            "scaling_type": "step_scaling",
            "step_adjustments": [
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 75,
                    "upper_bound": 85,
                    "adjustment": 1
                },
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 85,
                    "upper_bound": 95,
                    "adjustment": 3
                },
                {
                    "metric_name": "CPUUtilization",
                    "lower_bound": 95,
                    "upper_bound": 100,
                    "adjustment": 5
                }
            ],
            "evaluation_periods": 3,
            "metric_aggregation_type": "Average"
        }
    ]
}
```

(continues on next page)

(continued from previous page)

```

"instance_warmup": 300,
"scaling_metric": {
    "metric_name": "CPUUtilization",
    "namespace": "AWS/EC2",
    "comparison_operator": "GreaterThanOrEqualToThreshold",
    "evaluation_periods": 3,
    "evaluation_period": 60,
    "threshold": 75,
    "statistic": "Average",
    "dimensions": [
        {
            "name": "AutoScalingGroupName",
            "value": "$self"
        }
    ],
    "steps": [
        {
            "scalingAdjustment": 1,
            "metricIntervalUpperBound": 10,
            "metricIntervalLowerBound": 0
        },
        {
            "scalingAdjustment": 3,
            "metricIntervalUpperBound": 20,
            "metricIntervalLowerBound": 10
        },
        {
            "scalingAdjustment": 5,
            "metricIntervalLowerBound": 20
        }
    ],
    "disable_scale_in": false
}
]
}

```

`custom_scaling_policies` *Multiple Scaling Policies Example*

Note: An important consideration when scaling up, is also scaling down. This policy block scales up a cluster when CPUUtilization is GreaterThanThreshold of 50% by adding 5 servers. This policy block also scales down a cluster when CPUUtilization is LessThanThreshold of 25% and removes 1 server.

```

{
    "custom_scaling_policies": [
        {
            "scaling_type": "step_scaling",
            "instance_warmup": 300,
            "scaling_metric": {
                "metric_name": "CPUUtilization",
                "namespace": "AWS/EC2",

```

(continues on next page)

(continued from previous page)

```
"comparison_operator": "GreaterThanThreshold",
"evaluation_periods": 1,
"evaluation_period": 60,
"threshold": 50,
"statistic": "Average",
"dimensions": [
    {
        "name": "AutoScalingGroupName",
        "value": "$self"
    }
],
"steps": [
    {
        "scalingAdjustment": 5,
        "metricIntervalLowerBound": 0
    }
]
},
"disable_scale_in": false
},
{
"scaling_type": "step_scaling",
"instance_warmup": 300,
"scaling_metric": {
    "metric_name": "CPUUtilization",
    "namespace": "AWS/EC2",
    "comparison_operator": "LessThanThreshold",
    "evaluation_periods": 1,
    "evaluation_period": 300,
    "threshold": 25,
    "statistic": "Average",
    "dimensions": [
        {
            "name": "AutoScalingGroupName",
            "value": "$self"
        }
    ],
    "steps": [
        {
            "scalingAdjustment": -1,
            "metricIntervalUpperBound": 0
        }
    ]
},
"disable_scale_in": false
}
]
```

custom_scaling_policies Target Tracking Predefined Metrics Example

Note:

This policy used predefined target tracking metrics to perform target tracking. Refer

to the official AWS Documentation (or Spinnaker GUI) for list of predefined metric spec of each service such as:

[EC2 Predefined Metric Spec](#)

```
{
    "custom_scaling_policies": [
        {
            "scaling_type": "target_tracking",
            "instance_warmup": 180,
            "target_value": 60,
            "scaling_metric": {
                "metric_type": "predefined",
                "metric_name": "ASGAverageCPUUtilization"
            },
            "disable_scale_in": false
        }
    ]
}
```

[custom_scaling_policies Target Tracking Custom Metrics Example](#)

Note:

This policy used predefined target tracking metrics to perform target tracking. Refer to the official AWS Documentation (or Spinnaker GUI) for list of predefined metric spec of each service such as:

[EC2 Predefined Metric Spec](#)

```
{
    "custom_scaling_policies": [
        {
            "scaling_type": "target_tracking",
            "instance_warmup": 180,
            "target_value": 60,
            "scaling_metric": {
                "metric_type": "predefined",
                "metric_name": "ASGAverageCPUUtilization"
            },
            "disable_scale_in": false
        }
    ]
}
```

[custom_scaling_policies Keys](#)

[scaling_type](#)

Scaling Policy Type to use

Type: string

Options:

- "step_scaling"

- "target_tracking"

instance_warmup

Time period to wait before collecting metrics from cluster

Type: int

Default: 300

Units: seconds

disable_scale_in

You can disable the scale-in portion of a target tracking scaling policy. This feature provides you with the flexibility to scale in your Auto Scaling group using a different method. For example, you can use a different scaling policy type for scale in while using a target tracking scaling policy for scale out.

Policy: "target_tracking"

Type: boolean

Default: false

scaling_metric Keys

metric_type

Note: Only used with "scaling_type" policies that are of type target_tracking.

Specify the type of metric to use for evaluation: "custom" or "predefined"

If leveraging predefined metrics, ensure the name of the metric is supported via AWS Predefined Metric Spec (or via Spinnaker UI) such as: [EC2 Predefined Metric Spec](#)

Type: string

Default: None

Example Options:

- "predefined"
- "custom"

namespace

A namespace within a given provider to search metrics within.

Type: string

Default: None

Example Options:

- "AWS/EC2"
- "AWS/SQS"

metric_name

A Metric Name to search a given provider for. If metric_type is "predefined" for target tracking, ensure the name of the metric is supported via AWS Predefined Metric Spec API such as EC2 Predefined Metric Spec

Type: string

Default: None

Example Options:

- "CPUUtilization"
- "NetworkIn"
- "NetworkOut"
- "DiskReadBytes"

statistic

Statistic to calculate at the period to determine if threshold was met

Type: string

Default: None

Example Options:

- "Average"
- "Maximum"
- "Minimum"
- "Sum"

evaluation_periods

Count of evaluation periods to evaluate/check metrics (i.e. Average of 3 evaluation_periods every 30 second evaluation_period)

Type: int

evaluation_period

Time period in between metrics evaluations

Type: int

Units: seconds

threshold

Metric value threshold to begin scaling activities on.

Type: int

Default: None

comparison_operator

Comparison operator to perform against threshold

Type: string

Default: None

Example Options:

- "GreaterThanThreshold"
- "LessThanThreshold"

dimensions

Metric limitations such as specific metric values (such as AutoScalingGroupName, SQSQueueName, etc)

If a dimension name "AutoScalingGroupName" is provided and the value is equal to "\$self" Foremast will autodefine the current deployment.

Type: List of Objects

Default: None

Example Options:

```
{  
    "dimensions": [  
        {  
            "name": "AutoScalingGroupName",  
            "value": "$self"  
        }  
    ]  
}
```

steps

Steps to perform based on baseline threshold

Type: List of Objects

Default: None

Example Options:

Simple Scaling Bound

Note: Negative scalingAdjustment denote a removal operation.

```
{  
    "steps": [  
        {  
            "scalingAdjustment": -1,  
            "metricIntervalUpperBound": 0  
        }  
    ]  
}
```

Complex Upper and Lower Bound Range Step

Note:

`metricIntervalUpperBound` and
`metricIntervalLowerBound` are compounded on the threshold

I.E. do a step bound of 60-75% with a base "threshold" of 50 would have a "metricIntervalLowerBound" of 0 ($50+10=60$) and "metricIntervalUpperBound" of 20 ($50+25=75$))

```
{
  "steps": [
    {
      "scalingAdjustment": 3,
      "metricIntervalUpperBound": 25,
      "metricIntervalLowerBound": 10
    }
  ]
}
```

Scheduled Actions

Scheduled Actions are useful for scaling clusters on time based events such as provisioning shadow capacity, preparing for large spikes (such as email campaigns, promotions, holidays/sales, etc) and also deprovisioning (post spike). These operations work based on simple CRON expressions, making them easy to implement.

Note: Scheduled Actions persist between clusters as they are done at the service level. As a result, ensure you delete scheduled actions manually via the Spinnaker UI if you remove them from Foremast configuration files.

`scheduled_actions`

Defines scheduled actions to perform on an cluster group. You can specify multiple actions if needed. If this block does not exist, no scheduled actions will be attached.

`scheduled_actions Example`

```
{
  "scheduled_actions": [
    {
      "recurrence": "1 * * * *",
      "minSize": 1,
      "maxSize": 1,
      "desiredCapacity": 1
    },
    {
      "recurrence": "2 * * * *",
      "minSize": 1,
      "maxSize": 1,
      "desiredCapacity": 1
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
"minSize": 2,  
"maxSize": 2,  
"desiredCapacity": 2  
}  
]  
}
```

scheduled_actions Keys

recurrence

A valid CRON expression evaluated that is evaluated in UTC.

Type: string

minSize

Minimum size of Auto Scaling Group

Type: int

maxSize

Max size of Auto Scaling Group

Type: int

desiredCapacity

Desired Capacity of Auto Scaling Group

Type: int

Foremast Provider Tags

Foremast has the ability to perform a lot of important infrastructure actions, but Foremast is currently not stateful. This can cause issues with certain provider APIs that require some state (such as AWS S3 PutBucketNotification). In addition, some users may wish to restrict Foremast from making changes on specific resources. To address this, Foremast can leverage Tags/Labels to restrict some operations.

AWS Tags

This section addresses features that can be enabled/disabled by specific tags by AWS resources.

AWS S3 Bucket Tags

Foremast can leverage S3 bucket tags to allow/restrict specific Foremast features.

FOREMAST_LAMBDA_RESTRICT_BUCKET_NOTIFICATIONS

Restricts Foremast Lambda Deployment Triggers from overriding S3 bucket triggers

Type: string

Required: False

Example: "true" "false"

GCP Service Account IAM Policies

1.4 Advance Usages

- *Environment Variables*
- *Pipeline Configs*
- *Running Foremast*
 - *Method 1*
 - *Method 2*
- *Next Steps*

This section will show many advance usages of Foremast.

1.4.1 Environment Variables

These are environment variables used when executing Foremast

TRIGGER_JOB: The name of the Jenkins job that Spinnaker should look for as a trigger

APPNAME: The full name of your application in Spinnaker. \${GIT_REPO}\${PROJECT} is default

EMAIL: Email address associated with application in Spinnaker

PROJECT: The namespace or group of the application being set up

GIT_REPO: The name of the repo in the above namespace/group

RUNWAY_DIR: Path to the pipeline.json and application-master-\$account.json files created above

1.4.2 Pipeline Configs

The *pipeline.json* and *application-master-\$account.json* are critical files that determine on how each application in the pipeline will work. We recommend keeping these files in the same repository as your application but as long as they are on the same local machine as the Foremast runner they can be used.

In ~/runway create a file pipeline.json with the contents:

```
{  
    "deployment": "spinnaker",  
    "env": [ "account1", "account2"]  
}
```

In the same `~/runway` directory, create a file `application-master-$account.json` where `$account` is the same name as an account in your AWS credentials file and in your `env` list in `pipeline.json`. This file can be empty and it will just use the defaults provided at `application-master-$account.json`.

Note: You will need an `application-master-$account.json` config for each `$account` that you are deploying to.

See `pipeline.json` and `application-master-$account.json` for all configuration options.

1.4.3 Running Foremast

After setting up all of the configs there are a couple of ways to run Foremast components. You can use our bundled CLI endpoints that look at environment variables, or you can call each individual component with appropriate arguments on the CLI

Both methods will generate the same outcome. An application created in Spinnaker and a pipeline generated based on the configs.

Method 1

This is the recommended method on how to run Foremast. You need to first set the environment variables from above.

With the environment variables defined, you can simply run the command `foremast-pipeline` from the command line. This will create the Application in Spinnaker as well as generate a base pipeline.

Method 2

This method is more explicit and requires calling multiple Foremast components to create the configs, create the application, and generate the pipeline:

```
create-configs -o ./raw.properties -g ${PROJECT}/${GIT_REPO} -r ${RUNWAY_DIR}  
  
create-app -a ${APPNAME} --email ${EMAIL} --project ${PROJECT} --repo ${GIT_REPO}  
  
create-pipeline -a ${APPNAME} --triggerjob ${TRIGGER_JOB}
```

1.4.4 Next Steps

Take a look at the `Infrastructure` docs for details on the necessary Jenkins jobs.

Since Foremast is based on Jinja templates, it can be extended to do anything you need. For example, you can add functionality for creating scaling policies, setting up AWS infrastructure (elbs, security groups, iam policies, s3 buckets), sending slack notifications, and destroying old infrastructure. Take a look at our internal workflow docs for more detail on how Foremast is used at various organizations!

1.5 Infrastructure

- *Spinnaker*
 - *Spinnaker Component Versions*
- *AWS*
 - *AWS VPC Subnet Tags/Names*
 - *Foremast IAM Infrastructure*
 - *Foremast IAM Policy*
- *Jenkins*
 - *Necessary Jenkins Jobs*
- *Gitlab*

1.5.1 Spinnaker

- Foremast assumes that Spinnaker is already setup. Please see the [Spinnaker documentation](#) for assistance
- Requires connectivity to the Gate component of Spinnaker. Foremast also supports x509 authentication on Gate.
- Assumes AWS EBS is used for Packer bakes in Spinnaker Rosco

Spinnaker Component Versions

Previously, we used to publish internally tested versions. We have since moved to leveraging the release cadence set forth by the Spinnaker community. For more info around the Spinnaker release cadence, refer to the official can be found [Spinnaker Release Cadence](#) page.

For the latest releases of Spinnaker, check out the official [Spinnaker Release Versions](#) page.

If you have any issues with Foremast on the latest Spinnaker version, please file an issue (or pull request).

1.5.2 AWS

Foremast only works with AWS (for now). Below are the AWS requirements:

AWS VPC Subnet Tags/Names

Note: This is a general Spinnaker requirement when working with VPCs

- If new subnets are being setup, follow the [Spinnaker AWS Setup](#) guide.
- If using existing subnets add an `immutable_metadata` tag.
 - Example `immutable_metadata` tag: `{"purpose": "external", "target": "elb"}`
 - The "purpose" key will dictate how this appears in Spinnaker.

* Needs to be "internal" or "external" in order to properly work with Foremast

Foremast IAM Infrastructure

- A general IAM user/role will be needed for Foremast to work. In addition, Foremast will need credentials set up in a Boto3 configuration file. See [AWS Credentials](#) for details.
- Spinnaker handles the updates for things such as ELBs and security groups.

Foremast IAM Policy

Warning: The IAM Policy found below is a very generic policy for generic usage. You can and **should** consider locking down further using specific resource policies!

```
{  
    "Version": "2012-10-17",  
    "Statement": [  
        {  
            "Sid": "S3",  
            "Effect": "Allow",  
            "Action": [  
                "s3:PutObject",  
                "s3:PutLifecycleConfiguration",  
                "s3:PutEncryptionConfiguration",  
                "s3:PutBucketWebsite",  
                "s3:PutBucketVersioning",  
                "s3:PutBucketTagging",  
                "s3:PutBucketPolicy",  
                "s3:PutBucketNotification",  
                "s3:PutBucketLogging",  
                "s3:PutBucketCors",  
                "s3>ListObjects",  
                "s3>ListBucket",  
                "s3:GetObject",  
                "s3:GetBucketTagging",  
                "s3:GetBucketLocation",  
                "s3>DeleteObject",  
                "s3>DeleteBucketWebsite",  
                "s3>DeleteBucketPolicy",  
                "s3>CreateBucket"  
            ],  
            "Resource": "*"  
        },  
        {  
            "Sid": "EMR",  
            "Effect": "Allow",  
            "Action": "elasticmapreduce:*",  
            "Resource": "*"  
        },  
        {  
            "Sid": "Firehose",  
            "Effect": "Allow",  
            "Action": "firehose:*"  
        }  
    ]  
}
```

(continues on next page)

(continued from previous page)

```

    "Resource": "*"
},
{
  "Sid": "SQS",
  "Effect": "Allow",
  "Action": "sns:*",
  "Resource": "*"
},
{
  "Sid": "Kinesis",
  "Effect": "Allow",
  "Action": "kinesis:*",
  "Resource": "*"
},
{
  "Sid": "CloudwatchEvents",
  "Effect": "Allow",
  "Action": [
    "events:RemoveTargets",
    "events:PutTargets",
    "events:PutRule",
    "events>ListRules"
  ],
  "Resource": "*"
},
{
  "Sid": "Lambda",
  "Effect": "Allow",
  "Action": "lambda:*",
  "Resource": "*"
},
{
  "Sid": "APIGateway",
  "Effect": "Allow",
  "Action": "apigateway:*",
  "Resource": "*"
},
{
  "Sid": "DataPipeline",
  "Effect": "Allow",
  "Action": "datapipeline:*",
  "Resource": "*"
},
{
  "Sid": "SNS",
  "Effect": "Allow",
  "Action": [
    "sns:Unsubscribe",
    "sns:Subscribe",
    "sns>ListTopics"
  ],
  "Resource": "*"
},
{
  "Sid": "CloudwatchLogs",
  "Effect": "Allow",
  "Action": [

```

(continues on next page)

(continued from previous page)

```
        "logs:PutSubscriptionFilter",
        "logs:DescribeSubscriptionFilters",
        "logs>DeleteSubscriptionFilter"
    ],
    "Resource": "*"
},
{
    "Sid": "Route53",
    "Effect": "Allow",
    "Action": "route53:*",
    "Resource": "*"
},
{
    "Sid": "ELB",
    "Effect": "Allow",
    "Action": [
        "elasticloadbalancing:SetLoadBalancerPoliciesOfListener",
        "elasticloadbalancing:ModifyLoadBalancerAttributes",
        "elasticloadbalancing:DescribeLoadBalancers",
        "elasticloadbalancing>CreateLBCookieStickinessPolicy",
        "elasticloadbalancing>CreateAppCookieStickinessPolicy"
    ],
    "Resource": "*"
},
{
    "Sid": "EC2",
    "Effect": "Allow",
    "Action": [
        "ec2:RevokeSecurityGroupIngress",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:DescribeVpcs*",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups",
        "ec2:DeleteSecurityGroup",
        "ec2>CreateTags",
        "ec2>CreateSecurityGroup",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:AuthorizeSecurityGroupEgress"
    ],
    "Resource": "*"
},
{
    "Sid": "ASG",
    "Effect": "Allow",
    "Action": [
        "autoscaling:PutScheduledUpdateGroupAction",
        "autoscaling:DescribeScheduledActions",
        "autoscaling:DescribeAutoScalingGroups",
        "autoscaling>DeleteScheduledAction"
    ],
    "Resource": "*"
},
{
    "Sid": "IAM",
    "Effect": "Allow",
    "Action": [
        "iam:RemoveRoleFromInstanceProfile",
        "iam:ListInstanceProfilesForRole",
        "iam:ListAttachedRolePolicies",
        "iam:ListPoliciesForRole",
        "iam:ListRolePolicies",
        "iam:PutRolePolicy",
        "iam:DeleteRolePolicy",
        "iam:CreateRole",
        "iam:DeleteRole",
        "iam:UpdateRoleDescription",
        "iam:ListRoles"
    ]
}
```

(continues on next page)

(continued from previous page)

```

    "iam:PutRolePolicy",
    "iam:PassRole",
    "iam>ListServerCertificates",
    "iam>ListInstanceProfilesForRole",
    "iam:GetRole",
    "iam>CreateUser",
    "iam>CreateRole",
    "iam>CreateInstanceProfile",
    "iam>CreateGroup",
    "iam>AddUserToGroup",
    "iam>AddRoleToInstanceProfile"
],
"Resource": "*"
}
]
}

```

1.5.3 Jenkins

Foremast takes advantage of the Spinnaker Jenkins stage. In order for the Foremast generated pipeline to work you will need the following:

- Jenkins configuration named “JenkinsCI” in Spinnaker Igor
 - Example Igor config:

```

jenkins:
  Masters:
  -
    name: 'JenkinsCI' # The display name for this server
    address: 'http://jenkinsci.example.com'
    username: 'spinnaker'
    password: 'password'

```

Necessary Jenkins Jobs

The default generated pipeline requires a couple of Jenkins jobs to be setup in order to run.

- pipes-pipeline-prepare
 - Runs Foremast `prepare-infrastructure` during the “Infrastructure Setup” pipeline stage
 - Requires the following string variables
 - * PROJECT
 - * GIT_REPO
 - * ENV
 - * REGION
 - Example Shell after cloning Foremast:

```

virtualenv -p python3 venv
. venv/bin/activate
pip install -U --quiet .

```

(continues on next page)

(continued from previous page)

```
prepare-infrastructure
```

- pipes-scaling-policy
 - Runs Foremast `create-scaling-policy` for attaching a scaling policy if defined.
 - Only necessary if you plan on attaching scaling policies
 - Requires the following string variables
 - * PROJECT
 - * GIT_REPO
 - * ENV
 - * REGION
 - Example Shell after cloning Foremast

```
virtualenv -p python3 venv
. venv/bin/activate
pip install -U --quiet .

create-scaling-policy

# You can export these variables or also pass them beforehand such as:
export GIT_REPO=<repo_name>
export ENV=<spinnaker_env_name>

PROJECT=<repo_project> RUNWAY_DIR=<OS_path_to_runway_dir> \
REGION=<spinnaker_env_region> \
foremast-infrastructure
```

1.5.4 Gitlab

Gitlab is not required for Spinnaker but if it is already part of your infrastructure you can have Foremast directly look up the `pipeline.json` and `application-master-$account.json` files. You will need to get the Gitlab Token of a user that has permissions to the desired repository and set them in your `foremast.cfg / config.py`.

1.6 Foremast - AWS Lambda Pipelines

1.6.1 Lambda Triggers and Events

Foremast supports multiple Lambda events. These are configured in the `application-master-$account.json` config and set as a list under the `lambda_triggers` key.

- *Example Configuration*
- *Configuration Details*
 - *type*

- *api-gateway Trigger Keys*
 - * *api_name*
 - * *resource*
 - * *method*
 - * *api_type*
- *cloudwatch-event Event Pattern Trigger Keys*
 - * *rule_name*
 - * *rule_type*
 - * *rule_description*
 - * *event_pattern*
- *cloudwatch-event Schedule Trigger Keys*
 - * *rule_name*
 - * *rule_type*
 - * *rule_description*
 - * *schedule*
- *cloudwatch-logs Trigger Keys*
 - * *log_group*
 - * *filter_name*
 - * *filter_pattern*
- *dynamodb-stream Trigger Keys*
 - * *stream_arn*
 - * *table_arn*
 - * *batch_size*
 - * *batch_window*
 - * *parallelization_factor*
 - * *starting_position*
 - * *starting_position_timestamp*
 - * *max_retry*
 - * *split_on_error*
 - * *destination_config*
 - * *max_record_age*
- *kinesis-stream Trigger Keys*
 - * *stream_arn*
 - * *batch_size*
 - * *batch_window*

```
* starting_position
* max_retry
* split_on_error
* destination_config
* max_record_age
- s3 Trigger Keys
  * bucket
  * events
  * prefix
  * suffix
- sns Trigger Keys
  * topic
- sqs Trigger Keys
  * queue_arn
  * batch_size
```

Example Configuration

This example would go in the `application-master-$account.json` configuration file.

```
{
  "lambda_triggers": [
    {
      "type": "api-gateway",
      "api_name": "lambdatest-api",
      "resource": "/index",
      "method": "GET"
    },
    {
      "type": "cloudwatch-event",
      "rule_name": "app cron - 5min",
      "rule_type": "schedule",
      "rule_description": "triggers lambda function every five minutes",
      "schedule": "rate(5 minutes)"
    },
    {
      "type": "cloudwatch-event",
      "rule_name": "GuardDutyEvents",
      "rule_type": "event_pattern",
      "rule_description": "Trigger Lambda Function for every AWS GuardDutyEvent",
      "event_pattern": {"source": ["aws.guardduty"]}
    },
    {
      "type": "cloudwatch-logs",
      "log_group": "/aws/lambda/awslimit_test",
      "filter_name": "Trigger lambda on every WARNING message",
      "filter_pattern": ""
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

},
{
  "type": "dynamodb-stream",
  "table_arn": "arn:aws:dynamodb:us-east-1:111111111111:table/dynamotest-stream",
  "stream_arn": "",
  "batch_size": 100,
  "batch_window": 0,
  "starting_position": "TRIM_HORIZON",
  "max_retry": 3000,
  "split_on_error": true,
  "destination_config": {
    "OnFailure": {
      "Destination": "arn:aws:sns:us-east-1:111111111111:snstest-queue"
    }
  }
},
{
  "type": "kinesis-stream",
  "stream_arn": "arn:aws:kinesis:us-east-1:111111111111:stream/kinesistest-stream"
},
{
  "batch_size": 100,
  "batch_window": 0,
  "parallelization_factor": 1,
  "starting_position": "TRIM_HORIZON",
  "starting_position_timestamp": 1604617998,
  "split_on_error": true,
  "max_retry": 3000,
  "destination_config": {
    "OnFailure": {
      "Destination": "arn:aws:sqs:us-east-1:111111111111:sqstest-queue"
    }
  }
},
{
  "type": "s3",
  "bucket": "app-bucket-dev",
  "events": [
    "s3:ObjectCreated:*"
  ],
  "prefix": "",
  "suffix": ""
},
{
  "type": "sns",
  "topic": "app-dns-dev"
},
{
  "type": "sqs",
  "queue_arn": "arn:aws:sqs:us-east-1:111111111111:sqstest-queue",
  "batch_size": 10
}
]
}
}

```

Configuration Details

type

Specifies what type of Lambda event/trigger to use. This needs to be set for all events.

Type: string

Required: True

Options:

- "api-gateway" - API Gateway Lambda trigger
- "cloudwatch-event" - Cloudwatch Event Lambda trigger
- "cloudwatch-logs" - Cloudwatch Logs Lambda trigger
- "dynamodb-stream" - DynamoDB Stream Lambda trigger
- "kinesis-stream" - Kinesis Stream Lambda trigger
- "sns" - SNS Lambda trigger
- "sqS" - SQS Queue Lambda trigger
- "s3" - S3 Lambda trigger

api-gateway Trigger Keys

Sets up an API Gatway event to trigger a lambda function.

api_name

The name of an existing API Gateway. If not provided, an API will be created.

Type: string

Required: False

Default: {app_name}

resource

The API resource to tie the Lambda function to.

Type: string

Required: True

Example: "/test"

method

The API Method to trigger the Lambda function.

Type: string

Required: True

Example: "GET"

api_type

The API Type for the gateway integration.

Type: string

Required: False

Default: "AWS"

Values:

- "HTTP"
- "MOCK"
- "HTTP_PROXY"
- "AWS_PROXY"

cloudwatch-event Event Pattern Trigger Keys

A CloudWatch event pattern for Lambda triggers.

rule_name

The name of the CloudWatch rule being created.

Type: string

Required: True

rule_type

Type of CloudWatch Rule to create, must be set to "event_pattern" for Event Pattern Triggers.

Type: string

Required: True

Default: "schedule"

Values:

- "schedule"
- "event_pattern"

rule_description

Description of the rule being created.

Type: string

Required: False

event_pattern

CloudWatch Rule Event Pattern JSON. Usage Help can be found using the CloudWatch Rule GUI or the Docs: <https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/CloudWatchEventsandEventPatterns.html>

Type: string

Required: True

Examples:

- { "source": ["aws.guardduty"] }
- { "source": ["aws.ec2"], "detail-type": ["EC2 Instance State-change Notification"], "detail": { "state": ["running"] } }

cloudwatch-event Schedule Trigger Keys

A CloudWatch Scheduled event for Lambda triggers.

rule_name

The name of the CloudWatch rule being created.

Type: string

Required: True

rule_type

Type of CloudWatch Rule to create

Type: string

Required: False

Default: "schedule"

Values:

- "schedule"
- "event_pattern"

rule_description

Description of the rule being created.

Type: string

Required: False

schedule

The rate or cron string to trigger the Lambda function.

Type: string

Required: True

Examples:

- "rate(5 minutes)"
- "cron(0 17 ? * MON-FRI *)"

cloudwatch-logs Trigger Keys

A lambda event that triggers off a Cloudwatch log action.

log_group

The name of the log group to monitor.

Type: string
Required: True
Example: "/aws/lambda/test_function"

filter_name

The name of the filter on log event.

Type: string
Required: True

filter_pattern

The pattern to look for in the log_group for triggering a Lambda function.

Type: string
Required: True
Example: "warning"

dynamodb-stream Trigger Keys

A lambda event that triggers off a DynamoDB Stream.

Warning: Ensure IAM Role has permissions to the DynamoDB table/stream via "services" block

stream_arn

DynamoDB Stream ARN to use for triggering lambda.

Type: string
Required: True, if table_arn is not set.
Example: "arn:aws:dynamodb:us-east-1:111111111111:table/foremast-test/stream/2018-06-07T03:12:22.234"

table_arn

DynamoDB Table ARN to use for triggering lambda.

batch_size

The maximum number of items to retrieve in a single batch.

Type: int

Required: False

Default: 100

Max: 1000

batch_window

The maximum amount of time to gather records before invoking the function, in seconds.

Type: int

Required: False

Default: 0

Max: 300

parallelization_factor

For Kinesis Streams, the number of batches to process from each shard concurrently.

Type: int

Required: False

Default: 1

starting_position

The position in a stream from which to start reading.

Type: string

Required: False

Default: TRIM_HORIZON

Options:

- TRIM_HORIZON
- AT_TIMESTAMP - KINESIS STREAMS ONLY
- LATEST

starting_position_timestamp

The UTC timestamp (represented in Epoch Time) from which to start reading..

Type: int

Required: False

Default: None

max_retry

Skips retrying a batch of records when it has reached the Maximum Retry Attempts.

Type: int

Required: False

Default: 10000

Max: 10000

split_on_error

Breaks the impacted batch of records into two when a function returns an error, and retries them separately.

Type: boolean

Required: False

Default: false

Options:

- true
- false

destination_config

Continue processing records after error, and send metadata of bad data record to an SQS queue or SNS topic.

Type: string

Required: False

Default: “““

Options:

- arn:aws:sqs:us-east-1:111111111111:sqstest-queue
- arn:aws:sns:us-east-1:111111111111:snstest-queue

max_record_age

Maximum age of a record that is send to the function for processing.

Type: int

Required: False

Default: 604800

Max: 604800

kinesis-stream Trigger Keys

A lambda event that triggers off a Kinesis Stream.

Warning: Ensure IAM Role has permissions to the Kinesis Stream via "services" block

stream_arn

Kinesis Stream ARN to use for triggering lambda.

Type: string

Required: True

Example: "arn:aws:kinesis:us-east-1:111111111111:stream/kinesistest-stream"

batch_size

The maximum number of items to retrieve in a single batch.

Type: int

Required: False

Default: 100

Max: 10000

batch_window

The maximum amount of time to gather records before invoking the function, in seconds.

Type: int

Required: False

Default: 0

Max: 300

starting_position

The position in a stream from which to start reading.

Type: string

Required: False

Default: TRIM_HORIZON

Options:

- TRIM_HORIZON
- LATEST

max_retry

Skips retrying a batch of records when it has reached the Maximum Retry Attempts.

Type: int

Required: False

Default: 10000

Max: 10000

split_on_error

Breaks the impacted batch of records into two when a function returns an error, and retries them separately.

Type: boolean

Required: False

Default: false

Options:

- true
- false

destination_config

Continue processing records after error, and send metadata of bad data record to an SQS queue or SNS topic.

Type: string

Required: False

Default: “““

Options:

- arn:aws:sqs:us-east-1:111111111111:sqstest-queue
- arn:aws:sns:us-east-1:111111111111:snstest-queue

max_record_age

Maximum age of a record that is send to the function for processing.

Type: int

Required: False

Default: 604800

Max: 604800

s3 Trigger Keys

A Lambda trigger on S3 bucket actions.

bucket

The bucket of the event to monitor.

Type: string

Required: True

events

The S3 event to trigger the lambda function from.

Type: List

Required: True

Example: ["s3:ObjectCreated:*", "s3:ObjectRemoved:Delete"]

prefix

Sets up a prefix filter on S3 bucket events.

Required: False

Example: "logs/"

suffix

Sets up a suffix filter on s3 bucket events.

Required: False

Example: "jpg"

sns Trigger Keys

A Lambda trigger on SNS topic events.

topic

The SNS topic name to monitor for events.

Type: string

Required: True

sqs Trigger Keys

A Lambda trigger on SQS queue events.

queue_arn

SQS Queue ARN to use for triggering lambda.

Type: string

Required: True

Example: "arn:aws:sqs:us-east-1:111111111111:sqstest-queue"

batch_size

The maximum number of items to retrieve in a single batch.

Type: int

Required: False

Default: 10

Max: 10

1.6.2 Lambda Pipeline

- *Overview*
- *Lambda Specific Setup*
- *Lambda Pipeline Example*

Overview

Foremast supports the ability to setup Lambda infrastructure and build a pipeline around Lambda deployments. This was designed to be very similar to the default EC2 pipeline. It requires the same configuration files and general setup.

Lambda Specific Setup

1. Look at the *Getting Started* guide for basic setup. The Lambda process will be very similar
2. Look at the *lambda Block* configurations in *pipeline.json* and *application-master-\$account.json*.
3. In *pipeline.json* set "type" : "lambda" in order for Foremast to treat the application as a Lambda function.
4. Setup the desired Lambda triggers. See *Lambda Triggers and Events* for details.

Lambda Pipeline Example

1. Generate a ZIP artifact of your desired Lambda function
2. Trigger Spinnaker Lambda pipeline
3. Spinnaker runs “Infrastructure Setup Lambda”
 1. Sets up default function
 2. Sets up event triggers
 3. Sets up IAM Roles
 4. Sets up security groups
4. Spinnaker runs a “Deploy Lambda” stage
 1. This stage uploads the ZIP artifact to the created Lambda function
5. Manual Judgement checkpoint for deploying to the next environment
6. Repeat steps 3-5 for each desired environment



1.7 Foremast - AWS S3 Pipelines

1.7.1 S3 Pipeline

- *Overview*
- *S3 Specific Setup*
- *S3 Pipeline Example*

Overview

Foremast supports the ability to setup S3 infrastructure and build a pipeline around S3 deployments. This was designed to be very similar to the default EC2 pipeline. It requires the same configuration files and general setup.

S3 Specific Setup

1. Look at the [Getting Started](#) guide for basic setup. The S3 process will be very similar
2. Look at the [s3 Block](#) configurations in [pipeline.json](#) and [application-master-\\$account.json](#).
3. In [pipeline.json](#) set "type" : "s3" in order for Foremast to treat the application as an S3 deployment.

S3 Pipeline Example

1. Prepare a local folder containing your desired S3 deployment (commonly an uncompressed tar.gz)
2. Trigger Spinnaker S3 pipeline
3. Spinnaker runs “Infrastructure Setup S3”
 1. Sets up S3 bucket
 2. Attaches S3 bucket policies and metadata
 3. Creates friendly DNS record for s3 bucket if website enabled
4. Spinnaker runs a “Deploy S3” stage
 1. This stage uploads the local folder containing your artifacts to the created S3 bucket
5. Manual Judgement checkpoint for deploying to the next environment
6. Repeat steps 3-5 for each desired environment



1.8 Deploy Spinnaker Using Halyard

Run Spinnaker using Halyard to deploy Services to a Kubernetes Cluster. For more information about Spinnaker, see [Overview of Halyard Conventions](#).

1.8.1 Overview of Halyard Conventions

Quick commands:

```
# Deploy full Spinnaker in rolling fashion starting with bootstrap Services
hal deploy apply
```

Service Names

Type `list` in <https://github.com/spinnaker/halyard/blob/master/halyard-deploy/src/main/java/com/netflix/spinnaker/halyard/deploy/spinnaker/v1/service/SpinnakerService.java>.

- `clouddriver-bootstrap`
- **clouddriver**
 - Interfaces with all Cloud Providers: AWS, Kubernetes, etc.
- `consul-client`
- `consul-server`
- **deck**
 - Web UI served by Apache 2 by default
 - Talks directly to Gate for all information
- `echo`
- `flat`
- `front50`
- **gate**
 - Main entry point for every API call
 - The web UI makes calls to this API directly, which is why it needs to be publicly accessible
 - Makes backend calls to all Spinnaker Services
- **igor**
 - Interfaces with Continuous Integration (CI) Providers: Jenkins
 - Stores credentials for Git Repository Providers: GitHub
 - Scans for changes to trigger Pipelines
- **kayenta**
 - Service introduced into mainline version 1.7.0
 - Provides Automated Canary Analysis (ACA)
- `monitoring-daemon`
- `orca-bootstrap`

- orca
- redis-bootstrap
- redis
- **rosco**
 - Controls the Bake Stage for creating machine images: AMIs
 - Uses Packer underneath to provision machines and run configuration management
- vault-client
- vault-server

Service Settings

- <https://www.spinnaker.io/reference/halyard/custom/>

Override Kubernetes Service Settings

To override Kubernetes Service settings in the generated file `.hal/default/history/service-settings.yml`, create a file `.hal/default/service-settings/SERVICE.yml`.

Example `.hal/default/service-settings/echo.yml`:

```
kubernetes:  
  podAnnotations:  
    sumologic.com/format: text  
    sumologic.com/sourceCategory: spinnaker/echo  
    sumologic.com/sourceName: echo  
env:  
  JAVA_OPTS: -Xms2g -Xmx2g
```

Override Spring Profile Settings

To override settings Spring Profile settings, create a file `.hal/default/profiles/SERVICE-local.yml`.

Example `.hal/default/profiles/clouddriver-local.yml`:

```
serviceLimits:  
  cloudProviderOverrides:  
    aws:  
      rateLimit: 15
```

Swagger APIs

Most Spinnaker Services have a Swagger UI for exploration of the API hosted at <http://localhost:\protect\T1\textdollar\protect\T1\textbraceleftPORT\protect\T1\textbraceright/swagger-ui.html>. The only publicly facing Service with Swagger is Gate. Use Kubernetes to port forward for all other private Services.

```
kubectl --namespace spinnaker get pods      # Find the Pod for the Spinnaker Service  
kubectl --namespace spinnaker get services  # Find the exposed port  
kubectl --namespace spinnaker port-forward ${POD_NAME} ${PORT}  
# Go to http://localhost:\${PORT}/swagger-ui.html
```

Common Services with useful APIs:

- Clouddriver
- Gate

1.8.2 Problems

Issues that have come up when deploying or managing Spinnaker.

Contents

- *Problems*
 - *Kubernetes*
 - * *Pods in Unknown State*
 - *Fiat*
 - * *Fiat does not come up*
 - *Gate API SSL*
 - * *Gate not serving x.509 port*
 - * *Using a self-signed Certificate for Gate with Traefik Ingress controller*
 - * *Loading page shows 500 Internal Server Error*
 - * *Creating an Application will result in an Access denied error*
 - *Authorization*
 - * *Disable Clusters*
 - * *Traffic Guards*
 - *Provider Rate Limiting*
 - * *AWS throttling errors*
 - *Application Deployment*
 - * *Error when deploying an Application*
 - * *Exception (Determine Source Server Group) 403*
 - *Pipeline Trigger*
 - * *Pipelines not triggering when Fiat enabled*
 - *Memory Usage*
 - * *Microservices will grow and consume gratuitous amounts of RAM*
 - *Web UI*
 - * *Availability Zones do not show when creating a Load Balancer*
 - * *Create an internal load balancer not checked by default*

Kubernetes

Pods in Unknown State

- Seems to happen when **hal deploy apply** gives up after waiting on the bootstrap Services
- Not able to delete Pods
- Have to restart Docker Daemon on Nodes, or rotate Nodes out
- **Solution:**
 - Seems like this does not occur when running on Kubernetes Nodes with more resources available

Fiat

Fiat does not come up

- Shows error

```
2018-08-09 08:39:51.952 ERROR 1 --- [ecutionAction-6] c.n.s.fiat.roles.  
  ↳UserRolesSyncer           : [] Unable to resolve service account  
  ↳permissions.  
com.netflix.spinnaker.fiat.permissions.PermissionResolutionException: com.  
  ↳netflix.spinnaker.fiat.providers.ProviderException: (Provider:  
  ↳DefaultAccountProvider) retrofit.RetrofitError: connect timed out
```
- **Solution:**
 - Make sure Clouddriver has a Pod running
 - <https://github.com/spinnaker/fiat/blob/397706a98b56d4470a06f63972048a3157f98aaf/fiat-roles/src/main/java/com/netflix/spinnaker/fiat/providers/internal/ClouddriverService.java#L32-L36>
 - Make sure spec.replicas > 0

```
kubectl -n spinnaker get pods  
kubectl -n spinnaker get replicaset  
kubectl -n spinnaker edit replicaset spin-clouddriver-v##
```

Gate API SSL

Gate not serving x.509 port

- x.509 port defined as default.apiPort: 8085 in gate-local.yml
- Output of **netstat -ntlp** on Gate shows no listener on 8085
- **Solution:**

- Requires SSL to be enabled

```
hal config security api ssl enable
```

Using a self-signed Certificate for Gate with Traefik Ingress controller

- **hal config security api ssl enable**

- Loading page shows 502 Bad Gateway
- Traefik Ingress using HTTP to communicate with the new HTTPS port
- Traefik recognizes the scheme based on port, if 443 use HTTPS
- **Solution:**

- Configure Traefik to use HTTPS
- Update Gate Service with **kubectl** to route port 443

```
apiVersion: v1
kind: Service
metadata:
  name: spin-gate
  namespace: spinnaker
  annotations:
    prometheus.io/path: /prometheus_metrics
    prometheus.io/port: "8008"
    prometheus.io/scrape: "true"
spec:
  ports:
    - name: https
      port: 443
      targetPort: 8084
    - name: http
      port: 8084
      targetPort: 8084
```

- Update Gate Ingress to use Service port 443

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: spin-gate
  namespace: spinnaker
spec:
  rules:
    - host: gate.example.com
      http:
        paths:
          - path: /
            backend:
              serviceName: spin-gate
              servicePort: https
```

- Now page loads with 500 Internal Server Error

Loading page shows 500 Internal Server Error

- Traefik Ingress does not trust self-signed Certificate
- **Possible solutions:**
 - Use a publicly trusted Certificate
 - Add the private Certificate Authority to Traefik
 - Set `insecureSkipVerify = true` in Traefik's global configuration

- **Solution:**

- Short term, set `insecureSkipVerify = true`
- Add configuration file for Traefik

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: traefik-config
  namespace: kube-system
data:
  traefik.toml: |
    logLevel = "INFO"

    insecureSkipVerify = true
```

- Mount Traefik configuration file

```
kind: Deployment
apiVersion: extensions/v1beta1
metadata:
  name: traefik-ingress-controller
  namespace: kube-system
  labels:
    k8s-app: traefik-ingress-lb
spec:
  template:
    spec:
      containers:
        - image: traefik
          name: traefik-ingress-lb
          args:
            - --api
            - --kubernetes
          volumeMounts:
            - name: traefik-config
              mountPath: /etc/traefik
      volumes:
        - name: traefik-config
          configMap:
            name: traefik-config
```

- Page now loads as expected

Creating an Application will result in an Access denied error

- Front50 returns 403 (permission denied)
- Orca error in logs:

```
2018-05-29 14:14:59.937 ERROR 1 --- [ handlers-19] c.n.s.orca.q.
  ↪handler.RunTaskHandler : [] Error running UpsertApplicationTask
  ↪for orchestration[00000000-0000-0000-0000-000000000000]
retrofit.RetrofitError: 403
  at retrofit.RetrofitError.httpError(RetrofitError.java:40)
  at retrofit.RestAdapter$RestHandler.invokeRequest(RestAdapter.
  ↪java:388)
```

(continues on next page)

(continued from previous page)

```

at retrofit.RestAdapter$RestHandler.invoke(RestAdapter.java:240)
at com.sun.proxy.$Proxy106.get(Unknown Source)
at com.netflix.spinnaker.orca.front50.Front50Service$get.call(Unknown
Source)
at com.netflix.spinnaker.orca.front50.tasks.AbstractFront50Task.
fetchApplication(AbstractFront50Task.groovy:73)
at com.netflix.spinnaker.orca.applications.tasks.
UpsertApplicationTask.performRequest(UpsertApplicationTask.groovy:39)
at com.netflix.spinnaker.orca.applications.tasks.UpsertApplicationTask
$performRequest.callCurrent(Unknown Source)
at com.netflix.spinnaker.orca.front50.tasks.AbstractFront50Task.
execute(AbstractFront50Task.groovy:67)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$handle$1$1.
invoke(RunTaskHandler.kt:82)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$handle$1$1.
invoke(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.AuthenticationAwareKt$sam
$Callable$55f02348.call(AuthenticationAware.kt)
at com.netflix.spinnaker.security.AuthenticatedRequest.lambda
$propagate$1(AuthenticatedRequest.java:79)
at com.netflix.spinnaker.orca.q.handler.AuthenticationAware
$DefaultImpls.withAuth(AuthenticationAware.kt:49)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
withAuth(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$handle$1.
invoke(RunTaskHandler.kt:81)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$handle$1.
invoke(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$withTask$1.
invoke(RunTaskHandler.kt:173)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler$withTask$1.
invoke(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler$withTask$1.
invoke(OrcaMessageHandler.kt:47)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler$withTask$1.
invoke(OrcaMessageHandler.kt:31)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler$withStage
$1.invoke(OrcaMessageHandler.kt:57)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler$withStage
$1.invoke(OrcaMessageHandler.kt:31)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler
$DefaultImpls.withExecution(OrcaMessageHandler.kt:66)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
withExecution(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler
$DefaultImpls.withStage(OrcaMessageHandler.kt:53)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
withStage(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler
$DefaultImpls.withTask(OrcaMessageHandler.kt:40)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
withTask(RunTaskHandler.kt:51)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
withTask(RunTaskHandler.kt:166)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
handle(RunTaskHandler.kt:63)
at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.
handle(RunTaskHandler.kt:51)

```

(continues on next page)

(continued from previous page)

```
at com.netflix.spinnaker.q.MessageHandler$DefaultImpls.  
↳ invoke(MessageHandler.kt:36)  
    at com.netflix.spinnaker.orca.q.handler.OrcaMessageHandler  
    ↳ $DefaultImpls.invoke(OrcaMessageHandler.kt)  
        at com.netflix.spinnaker.orca.q.handler.RunTaskHandler.  
    ↳ invoke(RunTaskHandler.kt:51)  
        at com.netflix.spinnaker.orca.q.audit.  
    ↳ ExecutionTrackingMessageHandlerPostProcessor  
    ↳ $ExecutionTrackingMessageHandlerProxy.  
    ↳ invoke(ExecutionTrackingMessageHandlerPostProcessor.kt:47)  
        at com.netflix.spinnaker.q.QueueProcessor$pollOnce$1$1.  
    ↳ run(QueueProcessor.kt:74)  
        at java.util.concurrent.ThreadPoolExecutor.  
    ↳ runWorker(ThreadPoolExecutor.java:1149)  
        at java.util.concurrent.ThreadPoolExecutor$Worker.  
    ↳ run(ThreadPoolExecutor.java:624)  
        at java.lang.Thread.run(Thread.java:748)
```

- **Solution:**

- Set `fiat.cache.expiresAfterWriteSeconds: 0` in `fiat-local.yml` and `services.fiat.cache.expiresAfterWriteSeconds: 0` in `spinnaker-local.yml`
 - * <https://www.bountysource.com/issues/48656889-application-not-found-and-delay-issue-in-ui>
 - * Property needs to be set in both files
 - * Reduces the default 20 seconds

- **Application creation workflow now goes:**

- * Front50 responds 404 (not found) instead of 403 (access denied)

```
com.netflix.spinnaker.front50.exception.NotFoundException: Object  
↳ not found (key: exampleapplication)
```

- * Create Application
 - * Application exists immediately

Authorization

Disable Clusters

- Anyone is able to disable and enable Clusters
- Destroying a Cluster will disable the Cluster, then fail when destroying with error Access denied to account \${ACCOUNT}

- **Solution:**

- Will fail properly with Traffic Guards enabled for Cluster

Traffic Guards

- Anyone can modify the Traffic Guards for an Application

- After removing safety, someone can later disable a Cluster and take down traffic

Provider Rate Limiting

AWS throttling errors

- ThrottleException in Clouddriver logs

```
2018-05-09 01:36:48.681 INFO 1 --- [cutionAction-47] com.amazonaws.  
→latency : ServiceName=[AmazonElasticLoadBalancing],  
→ThrottleException=[com.amazonaws.services.elasticloadbalancingv2.model.  
→AmazonElasticLoadBalancingException: Rate exceeded (Service:  
→AmazonElasticLoadBalancing; Status Code: 400; Error Code: Throttling;  
→Request ID: 00000000-0000-0000-0000-000000000000)],  
→AWSErrorCode=[Throttling], StatusCode=[400, 200],  
→ServiceEndpoint=[https://elasticloadbalancing.us-west-2.amazonaws.com],  
→RequestType=[DescribeTargetHealthRequest], AWSRequestId=[00000000-0000-  
→0000-0000-000000000000, 00000000-0000-0000-000000000000],  
→HttpClientPoolPendingCount=0, RetryCapacityConsumed=0,  
→ThrottleException=1, HttpClientPoolAvailableCount=0, RequestCount=2,  
→HttpClientPoolLeasedCount=0, RetryPauseTime=[474.151],  
→RequestMarshallTime=[0.002], ResponseProcessingTime=[0.214],  
→ClientExecuteTime=[700.076], HttpClientSendRequestTime=[0.059, 0.048],  
→HttpRequestTime=[4.672, 42.883], RequestSigningTime=[0.082, 0.105],  
→CredentialsRequestTime=[0.002, 0.002, 0.003],  
→HttpClientReceiveResponseTime=[4.564, 27.471],
```

- Solution:

- Decrease allowed Provider API requests per second

- * <https://github.com/spinnaker/clouddriver/pull/1291>
- * <https://blog.armory.io/fine-grained-rate-limits-for-spinnaker-clouddriver/>

Application Deployment

Error when deploying an Application

```
Exception ( Monitor Deploy )
unable to resolve AMI imageId from ami-a5532fdd
```

- Solution:

- Fix where Clouddriver is trying to find AMIs
- Not sure what the **hal** command is, but modify `.hal/config` so `primaryAccount` is the Account to search

```
deploymentConfigurations:
- name: default
  providers:
    aws:
      primaryAccount: HALYARD_AWS_ACCOUNT_NAME
```

Exception (Determine Source Server Group) 403

```
Exception ( Determine Source Server Group )
403
```

- **Solution 1:**

- Missing READ permissions for Account
- Look at .hal/config for what Roles are listed under READ
- For Service Accounts, add the Role
- For Users, add the User to the Group in the SAML or other authentication Provider

- **Solution 2:**

- Deploy Stage application value does not match Spinnaker Application
- In the UI, the Cluster name should be the same as the Spinnaker Application

Pipeline Trigger

Pipelines not triggering when Fiat enabled

```
# Igor
2018-10-25 23:25:06.607 INFO 1 --- [RxIoScheduler-4] c.n.s.igor.jenkins.
↳ JenkinsBuildMonitor : [master=Jenkins:job=example-job] has no other builds
↳ between [Thu Oct 25 23:21:42 GMT 2018 - Thu Oct 25 23:24:00 GMT 2018], advancing
↳ cursor to 1540509840709

# Echo
2018-10-25 23:25:06.607 INFO 1 --- [IoScheduler-987] c.n.s.e.p.monitor.
↳ TriggerMonitor : Found matching pipeline example-application:example-
↳ pipeline
2018-10-25 23:25:06.607 INFO 1 --- [IoScheduler-987] c.n.s.e.p.orca.
↳ PipelineInitiator : Triggering Pipeline(example-application, example-
↳ pipeline, 00000000-0000-0000-000000000000) due to Trigger(00000000-0000-0000-
↳ 0000-000000000000, jenkins, Jenkins, example-job, null, gitlab, null, null, null,
↳ null, null, null, {}, null, {}, null, null, [], null, null, null, null,
↳ Pipeline(example-application, example-pipeline, 00000000-0000-0000-0000-
↳ 000000000000))
2018-10-25 23:25:06.608 INFO 1 --- [it-/orchestrate] c.n.s.e.p.orca.OrcaService
↳ : ---> HTTP POST http://spin-orca.spinnaker:8083/orchestrate
2018-10-25 23:25:06.651 INFO 1 --- [it-/orchestrate] c.n.s.e.p.orca.OrcaService
↳ : <--- HTTP 403 http://spin-orca.spinnaker:8083/orchestrate (45ms)
2018-10-25 23:25:06.693 ERROR 1 --- [ Retrofit-Idle] c.n.s.e.p.orca.
↳ PipelineInitiator : Retrying pipeline trigger, attempt 1/5
2018-10-25 23:25:27.023 ERROR 1 --- [ Retrofit-Idle] c.n.s.e.p.orca.
↳ PipelineInitiator : Error triggering pipeline: Pipeline(example-application,
↳ example-pipeline, 00000000-0000-0000-000000000000)

# Orca
2018-10-25 23:25:06.686 INFO 1 --- [0.0-8083-exec-8] c.n.s.o.c.OperationsController
↳ : [] received pipeline 00000000-0000-0000-000000000000:{...}
2018-10-25 23:25:06.687 INFO 1 --- [0.0-8083-exec-8] c.n.s.o.c.OperationsController
↳ : [] requested pipeline: {...}
```

(continues on next page)

(continued from previous page)

```
2018-10-25 23:25:06.687 INFO 1 --- [0.0-8083-exec-8] c.n.s.orca.front50.
↳Front50Service : [] ---> HTTP GET http://spin-front50.spinnaker:8080/
↳pipelines/example-application?refresh=false
2018-10-25 23:25:06.692 INFO 1 --- [0.0-8083-exec-8] c.n.s.orca.front50.
↳Front50Service : [] <--- HTTP 403 http://spin-front50.spinnaker:8080/
↳pipelines/example-application?refresh=false (5ms)
```

- **Solution:**

- Missing Run As User with Application READ and WRITE Permissions
- When not populated, the Run As User defaults to Anonymous
- When there are any Roles configured in the Application Permissions, Anonymous authorization no longer works
- Create a Service Account: <https://www.spinnaker.io/setup/security/authorization/service-accounts/>
- Configure Spinnaker Application Permissions to allow READ and WRITE for any Role the Service Account belongs to

Memory Usage

Microservices will grow and consume gratuitous amounts of RAM

- **Solution:**

- Set memory limits for Containers

- * <https://www.spinnaker.io/reference/halyard/component-sizing/>
- * Set Pod memory requests and limits in .hal/config

```
deploymentConfigurations:
- name: default
  deploymentEnvironment:
    customSizing:
      spin-cloudriver:
        limits:
          memory: 2Gi
```

- Set the JVM flags to be 80-90% .hal/default/service-settings/clouddriver.yml

```
env:
# 2GB * .8
JAVA_OPTS: -Xmx1638m
```

- –Xms should be 80-90% of Pod requests
- –Xmx should be 80-90% of Pod limits

Web UI

Availability Zones do not show when creating a Load Balancer

- JavaScript Console errors when selecting Account

```
TypeError: Cannot read property 'slice' of undefined
```

- **Solution:**

- Specify default Account and Region in Deck
- Use `.hal/default/profiles/settings-local.js` to override the defaults in `.hal/default/staging/settings.js`

```
window.spinnakerSettings.providers.aws.defaults = {  
  account: 'test',  
  region: 'us-east-5',  
  iamRole: 'DEFAULT_IAM_PROFILE',  
};
```

Create an internal load balancer not checked by default

- Have to remember to check *Create an internal load balancer* when creating Load Balancers

- **Solution:**

- Configure Deck to infer the Internal flag based on the Subnet Purpose name
- Use `.hal/default/profiles/settings-local.js` to override the defaults in `.hal/default/staging/settings.js`

```
window.spinnakerSettings.providers.aws.loadBalancers.  
  ↪inferInternalFlagFromSubnet = true;
```

1.8.3 Requirements

- Docker
- Kubernetes configuration file

1.8.4 Running

Launch the Halyard daemon and drop into a prompt with `launch_daemon.bash`. It mounts the Host directory `hal` without a dot prefix to `/home/spinnaker/.hal/` inside the Halyard Container. This is so the directory is more visible outside of the Container.

```
export KUBECONFIG=/fully/qualified/path/to/.kube/config  
./launch_daemon.bash
```

Show the configuration that will be deployed:

```
hal version list  
hal config  
hal config --help # Explore and set configurations  
hal deploy apply
```

Run the post deployment definition to clean up the bootstrap Pods with `post-deploy.yml`.

```
export KUBECONFIG=/fully/qualified/path/to/.kube/config  
kubectl apply --filename post-deploy.yml
```

1.8.5 Update Version

```
hal version list  
hal config version edit --version ${NEW_VERSION}  
hal deploy apply
```

1.9 How To Contribute

- *Getting Started*
 - *Commits*
 - *Branches*
 - *Documentation*
 - *Testing*
- *Code Submission*
 - *Code Improvement*
 - *Code Submission*
 - *Code Review*
 - *Code Acceptance*

Contributions to Foremast are welcome.

1.9.1 Getting Started

Commits

Follow [semantic commits](#) to make `git log` a little easier to follow.

chore something just needs to happen, e.g. versioning

docs documentation pages in `_docs/` or docstrings

feat new code in `src/`

fix code improvement in `src/`

refactor code movement in `src/`

style aesthetic changes

test test case modifications in `test/`

Examples commit messages:

- chore: v10.0
- docs: Add configuration setting
- feat: Create Lambda function
- fix: Retry upload on failure

- refactor: Extract duplicate code
- style: isort, YAPF
- test: Coverage around add permissions

Branches

Use [slash convention](#) with the same leaders as [Commits](#), e.g.:

- chore/v10.0
- docs/configs
- feat/lambda
- fix/deadlock
- refactor/debug_util
- style/lambda_whitespace
- test/lambda_permission

Documentation

- Use reStructuredText for docstrings and documentation
- For docstrings, follow [Example Google Style Python Docstrings](#)
- For documentation pages, follow the strong guidelines from Python with [Getting Started](#)

Note:

- Use `.rst` for regular pages
 - Use `.rest` for pages included using `.. include:: file.rest` (fixes a Sphinx issue that thinks references are duplicated)
-

Testing

Run any unit tests available in `./tests/`.

```
virtualenv venv
source ./venv/bin/activate
pip install -U -r requirements-dev.txt

tox
```

1.9.2 Code Submission

Code Improvement

1. See if an [Issue](#) exists
 - Comment with any added information to help the discussion

2. Create an [Issue](#) if needed

Code Submission

1. See if a [Pull Request](#) exists
 - Add some comments or review the code to help it along
 - Don't be afraid to comment when logic needs clarification
2. Create a Fork and open a [Pull Request](#) if needed

Code Review

- Anyone can review code
- Any [Pull Request](#) should be closed or merged within a week

Code Acceptance

Try to keep history as linear as possible using a *rebase* merge strategy.

1. One thumb up at minimum, two preferred
2. Request submitter to *rebase* and resolve all conflicts

```
# Update `master`
git checkout master
git pull

# Update `feat/new` Branch
git checkout feat/new
git rebase master

# Update remote Branch and Pull Request
git push -f
```

3. Merge the new feature

```
# Merge `feat/new` into `master`
git checkout master
git merge --ff-only feat/new
git push
```

4. Delete merged Branch

1.10 How To Create Releases

1.10.1 Creating a New Release

When releasing a new version, the following needs to occur:

1. Pull the latest main branch

```
git pull origin main
```

2. Ensure all test via tox pass
3. Add version Tag

```
git tag -a v#.##.#  
git push --tags
```

4. Github Actions won tag creation will build/publish to PyPI
5. Ensure proper build on: <https://test.pypi.org/project/foremast/#history>

1.11 src

1.11.1 src package

Subpackages

src.foremast package

Subpackages

src.foremast.app package

Submodules

src.foremast.app.spinnaker_app module

Module contents

src.foremast.autoscaling_policy package

Submodules

src.foremast.autoscaling_policy.create_policy module

Module contents

src.foremast.awslambda package

Subpackages

src.foremast.awslambda.api_gateway_event package

Submodules

[`src.foremast.awslambda.api_gateway_event.api_gateway_event module`](#)**Module contents**[`src.foremast.awslambda.cloudwatch_event package`](#)**Subpackages**[`src.foremast.awslambda.cloudwatch_event.destroy_cloudwatch_event package`](#)**Submodules**[`src.foremast.awslambda.cloudwatch_event.destroy_cloudwatch_event.destroy_cloudwatch_event module`](#)**Module contents****Submodules**[`src.foremast.awslambda.cloudwatch_event.cloudwatch_event module`](#)**Module contents**[`src.foremast.awslambda.cloudwatch_log_event package`](#)**Subpackages**[`src.foremast.awslambda.cloudwatch_log_event.destroy_cloudwatch_log_event package`](#)**Submodules**[`src.foremast.awslambda.cloudwatch_log_event.destroy_cloudwatch_log_event.destroy_cloudwatch_log_event module`](#)**Module contents****Submodules**[`src.foremast.awslambda.cloudwatch_log_event.cloudwatch_log_event module`](#)**Module contents**[`src.foremast.awslambda.event_source_mapping package`](#)**Submodules**

[src.foremast.awslambda.event_source_mapping.event_source_mapping module](#)

Module contents

[src.foremast.awslambda.s3_event package](#)

Subpackages

[src.foremast.awslambda.s3_event.destroy_s3_event package](#)

Submodules

[src.foremast.awslambda.s3_event.destroy_s3_event.destroy_s3_event module](#)

Module contents

Submodules

[src.foremast.awslambda.s3_event.s3_event module](#)

Module contents

[src.foremast.awslambda.sns_event package](#)

Subpackages

[src.foremast.awslambda.sns_event.destroy sns_event package](#)

Submodules

[src.foremast.awslambda.sns_event.destroy sns_event.destroy sns_event module](#)

Module contents

Submodules

[src.foremast.awslambda.sns_event.sns_event module](#)

Module contents

Submodules

[src.foremast.awslambda.awslambda module](#)

[src.foremast.awslambda.awslambdaevent module](#)

Module contents

[src.foremast.configs package](#)

Submodules

[src.foremast.configs.outputs module](#)

[src.foremast.configs.prepare_configs module](#)

Module contents

[src.foremast.datipeline package](#)

Submodules

[src.foremast.datipeline.datipeline module](#)

Module contents

[src.foremast.dns package](#)

Subpackages

[src.foremast.dns.destroy_dns package](#)

Submodules

[src.foremast.dns.destroy_dns.destroy_dns module](#)

Module contents

Submodules

[src.foremast.dns.create_dns module](#)

Module contents

[src.foremast.elb package](#)

Subpackages

[src.foremast.elb.destroy_elb package](#)

Submodules

[src.foremast.elb.destroy_elb.destroy_elb module](#)

Module contents

Submodules

[src.foremast.elb.create_elb module](#)

[src.foremast.elb.format_listeners module](#)

[src.foremast.elb.splay_health module](#)

Module contents

[src.foremast.gcp_iam package](#)

Submodules

[src.foremast.gcp_iam.create_iam_resources module](#)

[src.foremast.gcp_iam.policy module](#)

Module contents

[src.foremast.iam package](#)

Subpackages

[src.foremast.iam.destroy_iam package](#)

Submodules

[src.foremast.iam.destroy_iam.destroy_iam module](#)

Module contents

Submodules

[src.foremast.iam.construct_policy module](#)

[src.foremast.iam.create_iam module](#)

[src.foremast.iam.resource_action module](#)

Module contents

src.foremast.pipeline package**Submodules****src.foremast.pipeline.clean_pipelines module****src.foremast.pipeline.construct_pipeline_block module****src.foremast.pipeline.construct_pipeline_block_datapipeline module****src.foremast.pipeline.construct_pipeline_block_datapipeline module****src.foremast.pipeline.construct_pipeline_block_lambda module****src.foremast.pipeline.construct_pipeline_block_s3 module****src.foremast.pipeline.construct_pipeline_block_stepfunction module****src.foremast.pipeline.create_pipeline module****src.foremast.pipeline.create_pipeline_cloudfunction module****src.foremast.pipeline.create_pipeline_datapipeline module****src.foremast.pipeline.create_pipeline_lambda module****src.foremast.pipeline.create_pipeline_manual module****src.foremast.pipeline.create_pipeline_onetime module****src.foremast.pipeline.create_pipeline_s3 module****src.foremast.pipeline.create_pipeline_stepfunction module****src.foremast.pipeline.jinja_functions module****src.foremast.pipeline.reenumerate_stages module****Module contents****src.foremast.s3 package****Subpackages****src.foremast.s3.destroy_s3 package**

Submodules

`src.foremast.s3.destroy_s3.destroy_s3 module`

Module contents

Submodules

`src.foremast.s3.create_archaius module`

`src.foremast.s3.s3apps module`

`src.foremast.s3.s3deploy module`

Module contents

`src.foremast.scheduled_actions package`

Submodules

`src.foremast.scheduled_actions.create_scheduled_actions module`

Module contents

`src.foremast.securitygroup package`

Subpackages

`src.foremast.securitygroup.destroy_sg package`

Submodules

`src.foremast.securitygroup.destroy_sg.destroy_sg module`

Module contents

Submodules

`src.foremast.securitygroup.create_securitygroup module`

Module contents

`src.foremast.slacknotify package`

Submodules

src.foremast.slacknotify.slack_notification module**Module contents****src.foremast.stepfunction package****Submodules****src.foremast.stepfunction.stepfunction module****Module contents****src.foremast.utils package****Submodules****src.foremast.utils.apps module****src.foremast.utils.asg module****src.foremast.utils.awslambda module****src.foremast.utils.backoff module****src.foremast.utils.banners module****src.foremast.utils.credentials module****src.foremast.utils.deep_chain_map module****src.foremast.utils.dns module****src.foremast.utils.dynamodb_stream module****src.foremast.utils.elb module****src.foremast.utils.encoding module****src.foremast.utils.foremast_configs module****src.foremast.utils.gate module****src.foremast.utils.gcp_environment module****src.foremast.utils.generate_filename module**

[**src.foremast.utils.generate_s3_tags module**](#)

[**src.foremast.utils.get_cloudwatch_event_rule module**](#)

[**src.foremast.utils.get_sns_subscriptions module**](#)

[**src.foremast.utils.get_sns_topic_arn module**](#)

[**src.foremast.utils.google_iap module**](#)

[**src.foremast.utils.kayenta module**](#)

[**src.foremast.utils.lookups module**](#)

[**src.foremast.utils.pipelines module**](#)

[**src.foremast.utils.properties module**](#)

[**src.foremast.utils.roles module**](#)

[**src.foremast.utils.security_group module**](#)

[**src.foremast.utils.slack module**](#)

[**src.foremast.utils.subnets module**](#)

[**src.foremast.utils.tasks module**](#)

[**src.foremast.utils.templates module**](#)

[**src.foremast.utils.vpc module**](#)

[**src.foremast.utils.warn_user module**](#)

Module contents

Submodules

[**src.foremast.args module**](#)

[**src.foremast.consts module**](#)

[**src.foremast.destroyer module**](#)

[**src.foremast.exceptions module**](#)

[src.foremast.runner module](#)

[src.foremast.validate module](#)

[src.foremast.version module](#)

Module contents

Module contents

1.12 Changelog

1.12.1 v5.15.4 (2022-10-25)

- Chore(deps): bump pyjwt from 1.7.1 to 2.4.0. [dependabot[bot]]
Bumps [\[pyjwt\]\(https://github.com/jpadilla/pyjwt\)](#) from 1.7.1 to 2.4.0. - [\[Release notes\]\(https://github.com/jpadilla/pyjwt/releases\)](#) - [\[Changelog\]\(https://github.com/jpadilla/pyjwt/blob/master/CHANGELOG.rst\)](#) - [\[Commits\]\(https://github.com/jpadilla/pyjwt/compare/1.7.1...2.4.0\)](#)
 - updated-dependencies: - dependency-name: pyjwt
 - dependency-type: direct:production
 - ...

1.12.2 v5.15.3 (2021-11-30)

Fix

- Added missing arg and updated alias logic to work with update/create funcs. [Diego Nava]

1.12.3 v5.15.2 (2021-10-29)

Fix

- Exponential backoff for lambda resource update conflicts. [linjmeyer]

Other

- Chore: linting fixes. [linjmeyer]

1.12.4 v5.15.1 (2021-10-27)

Fix

- Lambda standalone zip code updates. [linjmeyer]
- Lambda package_type not exposed to foremast pipeline lambda stage. [linjmeyer]

1.12.5 v5.15.0 (2021-10-28)

Features

- Migrated to common foremast-utils (from legacy gogo-utils) [Joel Vasallo]
 - This will keep the project inline with the new open source standards set
 - Attributions remain, yet referring to contributions to project as Foremast Org (as its been 3-4 years now)
 - Migrated to local assets, still waiting and cleaning up gogo-utils to foremast-utils in PyPI.

At this time we will not be deleting the old packages, but in time they may no longer work. Foremast-Utils will continue to be maintained and in new versions of foremast be required (updated docs and build packages).

1.12.6 v5.14.2 (2021-10-22)

Fix

- Lambda tests and backwards compatibility. [linjmeyer]
- Fixing lambda test errors. [linjmeyer]

Features

- Backwards compatibility for lambda infra step. [linjmeyer]
- Added optional standalone foremast lambdas deploy CLI. [linjmeyer]
- Added infra support for Lambda docker images. [linjmeyer]

1.12.7 v5.14.1 (2021-10-18)

Fix

- Silently failing on ExpiredToken AWS calls. [linjmeyer]

Features

- Friendly error when env is not defined in pipeline json files. [linjmeyer]

Other

- Chore: linting fixes. [linjmeyer]

1.12.8 v5.14.0 (2021-09-22)

Fix

- Github Action upload to PyPI. [Joel Vasallo]
- Typo. [Joel Vasallo]

- Lint Issue. [Joel Vasallo]
- Missing key in tests. [Joel Vasallo]

Features

- Added support for specifying lambda subnet purpose to enabled public vpc enabled lambdas. [Joel Vasallo]

Other

- Chore: adding stale bot. [linjmeyer]

1.12.9 v5.13.0 (2021-09-08)

Fix

- Fixing conflicts in setup.py. [linjmeyer]
- Migrating to new Gitter. [Joel Vasallo]
- Added Github Actions Readme. [Joel Vasallo]
- File open linting issues in setup.py. [linjmeyer]

Features

- Ability to add webhooks to pipeline completion. [linjmeyer]

Other

- Merge branch ‘feature/completion-webhooks’ [Joel Vasallo]
- Merge branch ‘gitter-badger-gitter-badge-1’ [Joel Vasallo]
- Merge branch ‘gitter-badge-1’ of <https://github.com/gitter- badger/foremast> into gitter-badger-gitter-badge-1. [Joel Vasallo]
- Add Gitter badge. [The Gitter Badger]

1.12.10 v5.12.0 (2021-05-21)

Fix

- Attempting to only run tox tests on pushes to branch vs double. [Joel Vasallo]
- Removed pre-releases. [Joel Vasallo]
- Get_subnets spanning multiple lines. [Joel Vasallo]

Features

- Removed TravisCI, added releases, and updated releasing. [Joel Vasallo]
- Added noqa statement. [Joel Vasallo]
- Added support to pull version/tag. [Joel Vasallo]
- Adding test pushes and publish in Github actions. [Joel Vasallo]
- Adding Github Actions to run Tox on PRs. [Joel Vasallo]
- OSS Internal Quality Step from sample pipelines. [Joel Vasallo]

1.12.11 v5.11.0 (2021-03-15)

Fix

- More detailed example. [Joel Vasallo]
- Base template fixes. [Joel Vasallo]

Features

- Added docs for stepfunction support. [Joel Vasallo]
- Fixed linting. [Joel Vasallo]
- Added Update and Create Step Function Logic. [Joel Vasallo]
- Init Commit of StepFunction pipeline support. [Joel Vasallo]

Other

- Lint: fixes. [Joel Vasallo]
- Func: Added stepfunction to base types. [Joel Vasallo]

1.12.12 v5.10.0 (2021-03-12)

Fix

- Extra spacing in docs. [Joel Vasallo]
- Updating md fmt and readded legacy support for non-secrets manager pipelines. [Joel Vasallo]

Other

- Merge branch ‘feat-dataapi’ [Joel Vasallo]
- Merge branch ‘master’ into feat-dataapi. [Joel Vasallo]
- Removing extra comma. [Wayne Taylor]
- Accidentally updated rds-db section. Removing. [Wayne Taylor]
- Commit for data apis for RDS and Redshift. [Wayne Taylor]

- Removing extra comma. [Wayne Taylor]
- Accidentally updated rds-db section. Removing. [Wayne Taylor]
- Commit for data apis for RDS and Redshift. [Wayne Taylor]

1.12.13 v5.9.0 (2021-03-04)

Features

- Added custom_tags to tests. [Joel Vasallo]
- Upgraded foremast Travis to bionic - latest release for OpenSSL. [Joel Vasallo]
- Added new app level tags to s3 resources. [Joel Vasallo]
- Added support for custom tags on ec2 deploys. [Joel Vasallo]

Other

- Added support for lambda tags. [Joel Vasallo]

1.12.14 v5.8.2 (2021-03-02)

Fix

- Cryptography vulnerability patching. [Joel Vasallo]

1.12.15 v5.8.1 (2021-03-02)

Fix

- S3 canary/alpha deploy hotfix. [Joel Vasallo]

1.12.16 v5.8.0 (2021-01-27)

Fix

- AWS IAM put role policy no longer fails silently. [linjmeyer]
- Fixed an issue where setting LATEST or TRIM_HORIZON would fail due to timestamp default. [Joel Vasallo]

Other

- Chore: linting fix. [linjmeyer]
- Disabling complexity checks. [Joel Vasallo]

1.12.17 v5.7.0 (2020-12-09)

Features

- Added MediaConvert to Services Block. [Joel Vasallo]

1.12.18 v5.6.0 (2020-11-18)

Fix

- Dpl 1.10.16 rolled out and removing unneeded edge dependency. [Joel Vasallo]

Features

- Added function to generate pipeline id in manual templates. [linjmeyer]

Other

- Chore: Better code comments for generate_predictable_pipeline_id. [linjmeyer]

1.12.19 v5.5.2 (2020-11-13)

Fix

- Using Edge dpl to resolve issue with required importlib-metadata requirement. [Joel Vasallo]
- GCP foremast_groups label now double underscore seperated. [linjmeyer]
- Setup.py failure. [Joel Vasallo]
- Typo in templates schema name. [linjmeyer]

Features

- Manual templates now support default stage templates. [linjmeyer]
- Added starting_position_timestamp and parallelization_factor for Kinesis stream lambda triggers. [Joel Vasallo]

Other

- Chore(deps): bump cryptography from 2.5 to 3.2. [dependabot[bot]]
Bumps [cryptography](<https://github.com/pyca/cryptography>) from 2.5 to 3.2. - [Release notes](<https://github.com/pyca/cryptography/releases>) - [Changelog](<https://github.com/pyca/cryptography/blob/master/CHANGELOG.rst>) - [Commits](<https://github.com/pyca/cryptography/compare/2.5...3.2>)

1.12.20 v5.5.1 (2020-11-02)

Fix

- Missing comma in cloudfunctions example docs. [linjmeyer]
- Added Missing Licenses on contributions. [Joel Vasallo]
- Better error message when cloud function event resource has duplicate project id. [linjmeyer]

Features

- Pipeline notification support for all providers and options. [linjmeyer]

Other

- Chore: Added missing retry policy for cloud functions. [linjmeyer]
- Create codeql-analysis.yml. [Joel Vasallo]

1.12.21 v5.5.0 (2020-10-23)

Fix

- Added more approval_skip in core templates. [Joel Vasallo]
- Core Templates fixed to use approval_skip flag. [Joel Vasallo]

Features

- Added ability to skip approvals per environment with enforcement. [Joel Vasallo]
 - Set *approval_skip* in application-master-env.json in repo
 - Checks if Foremast Admins allow skip via config.py settings.
 - If approved by Foremast admins, and set by repo, set to value and skip in templates

1.12.22 v5.4.0 (2020-10-23)

Fix

- Matching format of cloudfunction timeout to api reqs. [Joel Vasallo]
- Docs issue with CloudFunction timeouts. [Joel Vasallo]

Features

- Added support to timeout approval stages after specified time. [Joel Vasallo]
 - Ability to maintain cleaner pipelines and fail faster than default 72 hour timeout

1.12.23 v5.3.0 (2020-10-08)

Features

- Allowing for provisioned throughput of lambda. [Wayne Taylor]

Other

- Fixing linting on awslambda.py. [Wayne Taylor]
- Test cases. [Wayne Taylor]

1.12.24 v5.2.9 (2020-10-07)

Fix

- Reverting change to runner PROVIDER env var. [linjmeyer]

Features

- IAM access control for Cloud Functions. [linjmeyer]

1.12.25 v5.2.8 (2020-09-23)

Fix

- Missing repo argument for GCP Svc Account IAM policy. [linjmeyer]
- 409 conflict handling for GCP svc account IAM policies. [linjmeyer]

1.12.26 v5.2.7 (2020-09-21)

Features

- Jinja template for default svc account IAM policy. [linjmeyer]

Other

- Docs: docs for gcp svc account iam policies. [linjmeyer]

1.12.27 v5.2.6 (2020-09-18)

Fix

- Existing svc accounts not being found in GCP IAM. [linjmeyer]

Other

- Chore: linting issues in gcp environment. [linjmeyer]

1.12.28 v5.2.5 (2020-09-16)

Fix

- Merge conflicts. [linjmeyer]
- Automatic retry on GCP svc account creation failure. [linjmeyer]
- KeyError when retrieving service accounts in a new project. [linjmeyer]

Other

- Merge branch ‘fix/gcp-no-serviceaccounts-key-error’ into master. [Joel Vasallo]

1.12.29 v5.2.4 (2020-09-08)

Fix

- Missing method error in aws lambda deploy. [linjmeyer]

1.12.30 v5.2.3 (2020-09-08)

Fix

- KeyError gcp_roles during GCP IAM step. [linjmeyer]

1.12.31 v5.2.2 (2020-09-04)

Fix

- Bug stripping too many template path chars. [linjmeyer]

1.12.32 v5.2.1 (2020-09-03)

Fix

- Default Pipeline Type if not specified is ec2. [Joel Vasallo]

1.12.33 v5.2.0 (2020-09-03)

Features

- Cloud function deploy now supports region option. [linjmeyer]
- Cloud Function Spinnaker pipeline support. [linjmeyer]
- Cloud Function support for event triggers. [linjmeyer]
- Cloud Function support for egress/ingress settings. [linjmeyer]
- GCP Cloud Functions regional VPC support. [linjmeyer]
- Added GCP Cloud Function deployments. [linjmeyer]

Other

- Chore: Testing TravisCI. [linjmeyer]
- Docs: Cloud Function VPC example is now more generic. [linjmeyer]
- Chore: Linting fixes. [linjmeyer]
- Docs: Updated GCP Cloud Function docs. [linjmeyer]
- Docs: Added GCP Cloud Function docs. [linjmeyer]

1.12.34 v5.1.0 (2020-09-03)

Fix

- Made infra subcommands help text more clear. [linjmeyer]

Features

- Foremast describe environments command. [linjmeyer]
- Print-environment is now a subcommand. [linjmeyer]
- Option to print GCP Environments and projects. [linjmeyer]

1.12.35 v5.0.0 (2020-08-20)

Fix

- Fixing manual_types duplicate check in config. [linjmeyer]

Features

- Pipeline type configurations now backwards compatible. [linjmeyer]
- Added GCP IAM group based permissions. [linjmeyer]
- Added support for GCP pub/sub roles. [linjmeyer]
- Added support for GCP Datastore IAM roles. [linjmeyer]

- GCP API caching, simplification of project role bindings. [linjmeyer]
- Support GCP Secrets Manager IAM. [linjmeyer]
- Support multiple gcp environments. [linjmeyer]

Other

- Chore: GCP infra code style changes. [linjmeyer]
- Chore: GCP infra code style change. [linjmeyer]
- Docs: Added GCP Infra and IAM docs. [linjmeyer]
- Chore: Using retry library in GCP IAM Api calls. [linjmeyer]
- Chore: Removed unused GCP IAM Code. [linjmeyer]
- Renamed project to project_prefix in gcp_roles. [linjmeyer]
- Fixing linting errors. [linjmeyer]
- Removed unneeded package in requirements.txt. [linjmeyer]
- Updated code comments in GCP IAM. [linjmeyer]
- Simplified GCP IAM into roles per project. [linjmeyer]
- Split foremast infra into GCP and AWS code paths. [linjmeyer]

1.12.36 v4.20.0 (2020-08-13)

Features

- Adding support for efs in lambda. [Wayne Taylor]

Other

- Linting on documentation. [Wayne Taylor]

1.12.37 v4.19.0 (2020-08-12)

- Fixed syntax and errors found by lint. [Diego Nava]
- Added support for lambda gateway integration type. [Diego Nava]

1.12.38 v4.18.1 (2020-07-16)

Fix

- Logging of Lambda Triggers for easier debugging. [Joel Vasallo]

1.12.39 v4.18.0 (2020-06-09)

Fix

- Template path for Parameter Store. [Joel Vasallo]

Features

- Added SSM ParameterStore Service. [Joel Vasallo]

1.12.40 v4.17.1 (2020-06-03)

Fix

- Linting issue. [Joel Vasallo]
- Wrong Config Lookup. [Joel Vasallo]
- Json lint. [Joel Vasallo]
- Tests missing lambda_subnet_count. [Joel Vasallo]

Features

- Lambda_subnet_count to limit subnets associated to VPC Enabled Lambdas. [Joel Vasallo]

1.12.41 v4.17.0 (2020-06-03)

- Update .travis.yml. [José Eduardo Montenegro Cavalcanti de Oliveira]

1.12.42 v4.16.0 (2020-05-22)

Features

- First commit of lambda destinations. [Wayne Taylor]

Other

- Fixing linting issue as per build. [Wayne Taylor]
- Update of lambda test cases. [Wayne Taylor]

1.12.43 v4.15.1 (2020-02-24)

Fix

- Reverting jinja function master commit. [linjmeyer]
- Removed infra step from Cloud Functions pipeline. [linjmeyer]

Features

- Added raise exception helper to Jinja functions. [linjmeyer]
- Added GCP cloud function generate pipeline. [linjmeyer]
- Added cloudfunctions to foremast pipeline. [linjmeyer]

Other

- Added function to raise exception in Jinja. [linjmeyer]
- Chore: Removed unused jinja2 file for cloud functions. [linjmeyer]
- Chore: fixed linting issues. [linjmeyer]
- Chore: Updated Cloud Function source code copyrights. [linjmeyer]

1.12.44 v4.15.0 (2020-02-12)

Fix

- Removed try-except to lower complexity factor. [Diego Nava]
- Added whitespaces where recommended. [Diego Nava]
- Updated long lines. [Diego Nava]

Other

- Added max_record_age parameter to lambda triggers. [Diego Nava]
- Lambda events: fixed sqs update, and added options to kinesis. [Diego Nava]

1.12.45 v4.14.0 (2020-02-05)

Fix

- Fixing app unit tests. [linjmeyer]
- Fixed provider in autoscaling policy. [linjmeyer]
- Fixed tests and linting. [linjmeyer]

Features

- Removed plugin system. [linjmeyer]

Other

- Chore: relative imports for foremast modules. [linjmeyer]

1.12.46 v4.13.1 (2020-01-13)

- Added S3 Pipeline to root index. [Joel Vasallo]

1.12.47 v4.13.0 (2020-01-13)

Fix

- Unit Test missing new deploy_strategy. [Joel Vasallo]
- Updated test for new mirror logic simplification. [Joel Vasallo]
- S3 test missing artifact_branch. [Joel Vasallo]

Features

- Added support for branch based S3 Deployments. [Joel Vasallo]
 - Allows mirroring of git branches to s3 buckets using versioned folders
 - refactor: renamed logic from uri to path for simplicity and understanding
 - refactor: removed dependency on mirror (bool) to deploy_strategy (str)

Other

- Updated S3 Pipeline Docs. [Joel Vasallo]
- Lint: Fixing lint issues. [Joel Vasallo]
- Docs: Update faulty lambda-tracing example JSON. [Mark Schleske]

1.12.48 v4.12.0 (2019-12-23)

Features

- Linting fixes on manual pipelines. [linjmeyer]
- One manual pipeline template can create multiple pipelines. [linjmeyer]

Other

- Docs: Added manual pipeline template info to docs. [linjmeyer]

1.12.49 v4.11.2 (2019-12-15)

- Update cloudwatch_event.py to reduce length of statement_id. [Neal K]

According to the boto3 logs, statement_id just needs to be something unique: https://boto3.amazonaws.com/v1/documentation/api/latest/reference/services/lambda.html#Lambda.Client.add_permission

When generating statementId, the end result in the current version is something like “foremast-XXXXXXX_cloudwatch_XXXXXXXXX_YYYYY” where XXXXXXXX is the app name and YYYY is the cloudwatch event name (see line 61 for where rule_name is set).

With this code change, you'll eliminate the first XXXXXXXXXX. This will reduce the occurrences of the following error for end users while keeping the statementId universally unique:

2019-12-14 23:47:38,953 [DEBUG] foremast.utils.awslambda:add_lambda_permissions:131 - Add permission error: An error occurred (ValidationException) when calling the AddPermission operation: 1 validation error detected: Value 'foremast-XXXXXXXXXX_cloudwatch_XXXXXXXXXXXXXX_YYYYYY' at 'statementId' failed to satisfy constraint: Member must have length less than or equal to 100

(note, in my case XXXXXX is rather long, but I'd come in under 100 chars if XXXXXXXXXX didn't appear twice).

1.12.50 v4.11.1 (2019-12-03)

Features

- Removed debugging j2 templates. [linjmeyer]
- Removed unused import in j2 function class. [linjmeyer]
- Removed unused normalize name j2 function. [linjmeyer]
- Added additional pipeline template helpers. [linjmeyer]

1.12.51 v4.11.0 (2019-11-27)

Features

- Add github token gate authentication. [Matias Puerta]

Other

- Fix lint issue. [Matias Puerta]

1.12.52 v4.10.2 (2019-11-21)

Fix

- Lint/doc issue. [Joel Vasallo]

Features

- Support for limiting lambda PutBucketNotification on tagged buckets. [Joel Vasallo]

Other

- Removed unsupported batch_window parameter from sqs event lambda trigger. [Diego Nava]

1.12.53 v4.10.1 (2019-11-07)

- Hotfix: SQS Event Source does not require starting position. [Joel Vasallo]
https://docs.aws.amazon.com/lambda/latest/dg/API_CreateEventSourceMapping.html

1.12.54 v4.10.0 (2019-10-25)

Fix

- Dict to str for event_pattern. [Joel Vasallo]
- Lint issues. [Joel Vasallo]
- Setup.py install targets. [Joel Vasallo]

Features

- Added support for AWS Lambda Event Pattern. [Joel Vasallo]

Other

- Defaulting to passing event json if no json_data passed. [Joel Vasallo]
- Added Python 3.8 support to tests. [Joel Vasallo]

1.12.55 v4.9.5 (2019-10-08)

Fix

- TemplateNotFound exceptions when manual pipelines use j2 include function. [linjmeyer]

Other

- Fixed linting. [linjmeyer]
- Removed local debugging files. [linjmeyer]
- Better logging/debug support for json decode errors in manual templates. [linjmeyer]
- Added logging for manual template jinja rendering exceptions. [linjmeyer]
- Fixed linting. [linjmeyer]
- Added better error handling. [linjmeyer]

1.12.56 v4.9.4 (2019-10-08)

Fix

- Index out of range. [Joel Vasallo]

1.12.57 v4.9.3 (2019-10-07)

- Fixed asg index indicator. [Diego Naval]

1.12.58 v4.9.2 (2019-10-07)

- Fixing lint issue with whitespace. [Wayne Taylor]
- Raise exception when no templates dir set. [Wayne Taylor]
- Data types for dlq and tracing are dict and not list. [Wayne Taylor]

1.12.59 v4.9.1 (2019-10-07)

- Adding docs for concurrency limits. [Wayne Taylor]

1.12.60 v4.9.0 (2019-10-07)

- Fixing rebase issues. [Wayne Taylor]
- Missing test case caused failure. [Wayne Taylor]
- Commit missed saves configs and tests. [Wayne Taylor]
- Adding lambda tracing capabilities. [Wayne Taylor]
- Changing block: [Wayne Taylor]
- Adding comments after review from Joel. [Wayne Taylor]
- Fixing linting issue in awslambda module. [Wayne Taylor]
- Fixing linting issue in awslambda module. [Wayne Taylor]
- Adding tests and fixes. [Wayne Taylor]
- Adding initial support for dlq. [Wayne Taylor]
- Style: Fix file isort. [Sijis Aviles]
- Style: isort non-init files. [Sijis Aviles]
- Chore: Add isort checks. [Sijis Aviles]

Ignoring exit code initially as we fix any warnings.

1.12.61 v4.8.3 (2019-10-04)

Fix

- Unneeded loop of accounts on app creation. [Joel Vasallo]

1.12.62 v4.8.2 (2019-10-05)

Fix

- Lint issues. [Joel Vasallo]

1.12.63 v4.8.1 (2019-10-04)

Fix

- Fixed issue where GATE_AUTHENTICATION not provided. [Joel Vasallo]

1.12.64 v4.8.0 (2019-10-04)

Features

- Adding support for Google IAP Auth. [Joel Vasallo]

Other

- Added docs to configuration files. [Joel Vasallo]
- Added conditional enable/disable to prevent issues as well as default example update. [Joel Vasallo]
- Added Exception in the event of failed token call. [Joel Vasallo]
- Fixed issues with bad header. [Joel Vasallo]
- Added Fix to Get Creds. [Joel Vasallo]

1.12.65 v4.7.0 (2019-10-04)

Fix

- Missing import. [Joel Vasallo]
- Additional named URIs. [Joel Vasallo]
- Missing uri. [Joel Vasallo]
- Index of of range. [Joel Vasallo]
- More mocks. [Joel Vasallo]
- Import at wrong level. [Joel Vasallo]

Features

- Ability to change VPC name when looking up VPC ID. [Joel Vasallo]

Other

- Removed from travis.yml. [Joel Vasallo]
- Refactor: Removed support for Python versions 3.4 and 3.5. [Joel Vasallo]
- Fixed issue with generate pipeline. [Joel Vasallo]
- Updated to remove unimplemented feature and upper the method. [Joel Vasallo]
- Fixed comment. [Joel Vasallo]
- Lint: New line at end of file. [Joel Vasallo]

- Lint: fixed lint issues. [Joel Vasallo]
- Fixed test cases. [Joel Vasallo]
- Fixed unneeded mocks and missing import. [Joel Vasallo]
- Mock credentials call. [Joel Vasallo]
- Fixed mocks. [Joel Vasallo]
- Testing mock fix. [Joel Vasallo]
- More Mock fixes. [Joel Vasallo]
- Removed unused mocks as they have moved. [Joel Vasallo]
- Fixing tests to mock new request. [Joel Vasallo]
- SG: Fixed format of URI. [Joel Vasallo]
- Migrated over to gate util method. [Joel Vasallo]
- Converted to singular gate requests. [Joel Vasallo]
- Refactor: Centralizing gate calls to a single function. [Joel Vasallo]
- Google_iap func now returns the entire token response. [linjmeyer]
- Linting fixes. [linjmeyer]
- Added util function for Google Identity Aware Proxy tokens. [linjmeyer]
- Added docs for FOREMAST_CONFIG_FILE. [linjmeyer]

1.12.66 v4.6.0 (2019-09-20)

- Added AttributeError exception handling to dynamic config load. [linjmeyer]
- Fixed bug causing incorrect config py file to be loaded when using FOREMAST_CONFIG_FILE env. [linjmeyer]

1.12.67 v4.5.0 (2019-09-16)

Fix

- Lint issue. [Joel Vasallo]
- Wrong source names in check. [Joel Vasallo]
- Simplified logic for checking create event sources. [Joel Vasallo]
- Conditional check on trigger type. [Joel Vasallo]
- Added ability to upsert existing function and fixed docs. [Joel Vasallo]
- Exceptions and logging. [Joel Vasallo]

Features

- Generic support for event_source events instead of just DynamoDB. [Joel Vasallo]
- Converting DynamoDB Logic to generic event source mapping. [Joel Vasallo]

Other

- Fixed lint. [Joel Vasallo]
- Fixed lint errors. [Joel Vasallo]
- Removed unused import. [Joel Vasallo]
- Removed unused destroy. [Joel Vasallo]
- Fixed additional references to streams. [Joel Vasallo]
- Fixed incorrect imports. [Joel Vasallo]
- Converted pluralized references to reflect a singular stream as expected. [Joel Vasallo]
- Fixes to docs and code style suggestions. [Joel Vasallo]
- Added Docs for DynamoDB Streams and fixed table logic. [Joel Vasallo]
- Fixed logic to properly attach DynamoDB stream to Lambda. [Joel Vasallo]
- Updated Trigger Type for DynamoDB-Streams. [Joel Vasallo]
- Fixed logic calling stream lookups. [Joel Vasallo]
- Fixed import statements. [Joel Vasallo]
- Fixed linting issues. [Joel Vasallo]
- Updated dynamodb_streams util functions. [Joel Vasallo]
- Updated module name to dynamodb_streams. [Joel Vasallo]
- Added code to get DynamoDB Streams ARN. [Joel Vasallo]
 - IF ARN Provided, return.
 - if table arn provided, parse and get latest streams arn.
 - if table name provided, get latest streams arn.
- Added utils to get DynamoDB Table ARN. [Joel Vasallo]
- Added proper function naming. [Joel Vasallo]
 - Code not feature complete
- WIP: Base commit to add DynamoDB support for lambda triggers. [Joel Vasallo]

1.12.68 v4.4.1 (2019-09-12)

Fix

- Error when getting task status. [Sijis Aviles]

Sending the content-type header seems to cause a 400 bad request. Removing this should not cause a problem as we are doing a GET request.

1.12.69 v4.4.0 (2019-09-09)

Fix

- Self to \$self and docs improvements. [Joel Vasallo]
- Lint issue with continuation line under-indented for visual indent. [Joel Vasallo]
- Scaling_policy check and Jinja2 errors. [Joel Vasallo]
- Logic check for custom scaling policies. [Joel Vasallo]

Other

- Docs: Added docs for manual_pipelines. [Joel Vasallo]
- Lint: fixing lint issues. [Joel Vasallo]
- Doc: fixed typo in json example. [Joel Vasallo]
- Docs: Added and cleaned up docs around scaling policies/scheduled actions. [Joel Vasallo]
- Added helper logic for ASG Name and base docs for scaling policy. [Joel Vasallo]
- Fixing lint issues. [Joel Vasallo]
- Fixed Custom Step Scaling Json. [Joel Vasallo]
- Cleaned up template. [Joel Vasallo]
- Fixed templates and cleaned up flow for step vs tracking. [Joel Vasallo]
- Setup hooks to insert custom scaling policies. [Joel Vasallo]
- Added Templates for Custom Scaling. [Joel Vasallo]

1.12.70 v4.3.0 (2019-09-09)

- Added year to copyright kayenta.py. [linjmeyer]
- Fixed linting. [linjmeyer]
- Exposed app_name, group_name, repo_name and trigger_job to manual j2 templates. [linjmeyer]
- Updated for PR comments. [linjmeyer]
- Moved get_canary_id function into new utility file kayenta.py. [linjmeyer]
- Fixed linting issues. [linjmeyer]
- Updated doc strings. [linjmeyer]
- Moved get_canary_id function into utilites, minor refactoring to jinja_functions. [linjmeyer]
- Template variable key in pipeline.json and within templates is now configurable. [linjmeyer]
- Added support for custom defined template:// schemas and overriding manual type, e.g. k8s or kubernetes. [linjmeyer]
- Added support for config defined TEMPLATES_SCHEME_IDENTIFIER constant. [linjmeyer]
- Added safe checks for j2 vs. json files and safely getting user defined variables. [linjmeyer]
- Added support for shared templates in foremast templates directory. [linjmeyer]

- Get_canary_id jinja function supports owning application name as optional param. [linjmeyer]
- Linting fixes 2. [linjmeyer]
- Linting fixes. [linjmeyer]
- Updated copyright to 2019, reordered imports for manual pipeline. [linjmeyer]
- Renamed manual pipeline json_dict variable to pipeline_dict. [linjmeyer]
- Manual variables now accessible via variables.var_name. Added ability to expose some functions to jinja templates. [linjmeyer]
- Manual pipeline now supports jinja2 templating. [linjmeyer]

1.12.71 v4.2.2 (2019-09-04)

Fix

- Missing s3path for bucket name deploys. [Joel Vasallo]

1.12.72 v4.2.1 (2019-08-30)

- Security: Fixed Default S3 Bucket Website Enablement. [Joel Vasallo]
- Adding bucket_name for deploy option. [Wayne Taylor]

1.12.73 v4.2.0 (2019-08-28)

Features

- Scheduled Actions for clusters. [Joel Vasallo]

Other

- Cleaned up additional lint issues. [Joel Vasallo]
- Fixing more lint issues. [Joel Vasallo]
- Fixed lint issues. [Joel Vasallo]

1.12.74 v4.1.0 (2019-08-28)

- Make changes per code review. [Wayne Taylor]
- Adding support to override s3 bucket name. [Wayne Taylor]

1.12.75 v4.0.1 (2019-08-12)

Fix

- Added fix on bucket notifications null check. [Joel Vasallo]
- QE docs issue. [Joel Vasallo]

Other

- Docs: Added support explaining the undocumented QE Block. [Joel Vasallo]
 - Some questions as to what this is used for. Added some example keys that could be used to achieve quality testing.
 - Highlight the fact that this is customizable per company

1.12.76 v4.0.0 (2019-08-08)

Fix

- Return useful data from App creation. [Nate Tangsurat]
- Remove unused import. [Sijis Aviles]
- Missing import. [Sijis Aviles]
- Adding provider and resource properties to app plugin. [Sijis Aviles]
- Rename parameter to more appropriate name. [Sijis Aviles]
- Use different import for better testing. [Sijis Aviles]
- Ensure object is persisted. [Sijis Aviles]
- Simplify call to plugin. [Sijis Aviles]

Features

- Add plugin manager wrapper for runner. [Sijis Aviles]
- Add provider environment variable. [Sijis Aviles]
- Add provider argument. [Sijis Aviles]
- Convert app entry point to work with plugin. [Sijis Aviles]
- Split app creation into a plugin. [Sijis Aviles]
- Add BasePlugin structure. [Sijis Aviles]
- Add plugin manager. [Sijis Aviles]

Other

- Test: Check handling duplicate key for instance links. [Sijis Aviles]
- Test: Combined instance links. [Sijis Aviles]
- Refactor: Provider more appropriate name for method. [Sijis Aviles]
- Refactor: Simplify rendering of application template. [Sijis Aviles]
- Refactor: Simplify merging of default and config instance links. [Sijis Aviles]
- Chore: Put debug output in better location. [Sijis Aviles]
- Refactor: Plugins must specify provider and resource properties. [Sijis Aviles]
- Test: Simplify instance links tests. [Sijis Aviles]

- Refactor: Move `__init__` to base class. [Sijis Aviles]
- Refactor: Simplify plugin manager. [Sijis Aviles]
- Refactor: Use plugin for creating app in runner. [Sijis Aviles]
- Refactor: Raise exception if a plugin is not found. [Sijis Aviles]
- Refactor: PluginManager path begins at package root. [Sijis Aviles]
- Refactor: Use provider argument when creating app. [Sijis Aviles]
- Chore: Disable abstract-method pylint warnings. [Sijis Aviles]
- Chore: Add docstring to plugin manager. [Sijis Aviles]
- Style: YAPF. [Sijis Aviles]
- Fixed logging unknown values before control flow. [Joel Vasallo]

1.12.77 v3.58.0 (2019-07-30)

Fix

- Added ‘manual’ pipeline type to default allowed types. [Lin Meyer]

Other

- Adding support to allow data pipeline in service block so lambda can have permissions to trigger. [Wayne Taylor]

1.12.78 v3.57.0 (2019-07-12)

- Removing Prospector Yaml. [Joel Vasallo]
- Moved options to prospector yaml. [Joel Vasallo]
- Testing Prospector Profile. [Joel Vasallo]
- Reverting to Prospector. [Joel Vasallo]
- PyLint with Xargs. [Joel Vasallo]
- Fixed PyLint and some recommendations. [Joel Vasallo]
- Converting to native pylint for linting. [Joel Vasallo]
- Testing updated tox interpreter. [Joel Vasallo]
- Added Lambda Layers documentation. [Joel Vasallo]
- Fixed Lambda Tests to account for new lambda layers block. [Joel Vasallo]
- Initial support for Lambda Layers. [Joel Vasallo]
- Fixed Lamdba doc issue. [Joel Vasallo]

1.12.79 v3.56.1 (2019-06-07)

Fix

- Changed logic around scaling policy to account for check. [Joel Vasallo]

1.12.80 v3.56.0 (2019-06-07)

Features

- Added support to increment/decrement by custom amounts in scaling policies. [Joel Vasallo]
- Added support for S3 Bucket Notification Configurations. [Joel Vasallo]

1.12.81 v3.55.0 (2019-04-23)

Fix

- Strip leading slash. [Sijis Aviles]
In some scenarios the generated path ends up being //path, instead of /path.

Other

- Refactor: Remove murl from creating urls. [Sijis Aviles]

Using murl strips out the /path part of a host. This becomes apparent under the following scenarios:

- Using a gate proxy
- Setting up gate under a subpath url

In these scenarios the host is specified as <https://gate.url/path>, then the calls to gate should be https://gate.url/path/api_endpoint, instead they become https://gate.url/api_endpoint.

This change allows those hosts definitions to work as expected.

- Chore: Remove Gate() [Sijis Aviles]

This is no longer used, so its being removed.

- Fixed lambda environment default. [dnaval13]

Fixes #333

1.12.82 v3.54.0 (2019-03-11)

- Added KMS support for Decrypt of encrypted keys. [Joel Vasallo]
- Adding support for AWS Secrets Manager. [Joel Vasallo]
- Added tests to check runway_base_path defaults properly. [Joel Vasallo]
- Added support for custom runway directory name. [Joel Vasallo]
- Removed extra parameter from userdata. [Joel Vasallo]
- Pass repo and group_name to trigger. [Joel Vasallo]

1.12.83 v3.53.0 (2019-02-14)

Fix

- Consts string to native enforcing string in ast.literal_eval. [Joel Vasallo]
- Issue with missing keys in vpc and subnet return. [Joel Vasallo]

Other

- Fixed subnet check and updated vpc url. [Joel Vasallo]
- Update pipeline_prod_s3.json.j2. [dnaval13]

1.12.84 v3.52.0 (2019-01-20)

Fix

- Deploy strategy list. [Nate Tangsurat]
- Code-block directive formatting. [Nate Tangsurat]

Other

- Added docs for archaius_enabled. [Joel Vasallo]
- Converted Archaius to feature toggle to remove assumption of internal logic. [Joel Vasallo]
 - Not everyone uses archaius in the community, let alone uses an application.properties. This at least converts this logic to a toggle until this can be investigated further in a future release.
- Docs: Switch to Sphinx provided Napoleon. [Nate Tangsurat]
- Docs: Problem when triggering without run as user. [Nate Tangsurat]

When Fiat is enabled and an Application is configured with any Permissions, the Pipeline Trigger must use a Service Account in the Run As User field.

 - Docs: Halyard overview. [Nate Tangsurat]
 - Docs: Problems from running Halyard. [Nate Tangsurat]

1.12.85 v3.51.7 (2018-09-06)

Fix

- Update Generator before rendering Pipelines. [Nate Tangsurat]

Other

- Test: Update to use an Object. [Nate Tangsurat]

1.12.86 v3.51.6 (2018-09-05)

Fix

- Less redundant error message. [Nate Tangsurat]

Features

- Expose Pipeline Type to User Data. [Nate Tangsurat]
Need a way for User Data to know about the Pipeline Type.

Other

- Style: New line after Returns Sections. [Nate Tangsurat]

1.12.87 v3.51.5 (2018-08-29)

Fix

- Mark setup executable. [Nate Tangsurat]

Features

- Passthrough variables for Pipeline blocks. [Nate Tangsurat]
See also: #320

1.12.88 v3.51.4 (2018-08-29)

Fix

- Newer linting issues. [Nate Tangsurat]
- Unpin prospector for testing. [Nate Tangsurat]
- Fixture decorator goes last. [Nate Tangsurat]

Features

- Environment variable FOREMAST_CONFIG_FILE. [Nate Tangsurat]
Support setting the dynamic configuration file to use using the environment variable *FOREMAST_CONFIG_FILE*.

1.12.89 v3.51.3 (2018-08-16)

- Adding support for passing custom task uri. [Joel Vasallo]

1.12.90 v3.51.2 (2018-08-01)

Fix

- Need to pass configuration format strings. [Nate Tangsurat]

1.12.91 v3.51.1 (2018-07-30)

Fix

- Need “location” in Traffic Guards. [Nate Tangsurat]

1.12.92 v3.51.0 (2018-07-26)

Fix

- Use formatted attributes from gogoutils. [Nate Tangsurat]
- Pass name Generator to User Data. [Nate Tangsurat]
- Pass Generator down to Pipeline setup. [Nate Tangsurat]

Other

- Refactor: Remove extra keyword argument. [Nate Tangsurat]
- Refactor: Unused import. [Nate Tangsurat]

1.12.93 v3.50.4 (2018-07-25)

Fix

- Pass *gogoutils.Generator* to create app template. [Sijis Aviles]
- Pass DEFAULT_RUN_AS_USER to create app template. [Sijis Aviles]

Other

- Test: Pass all parameters so Generator() works. [Sijis Aviles]
- Chore: Fix urls after org move. [Sijis Aviles]

1.12.94 v3.50.3 (2018-07-24)

Fix

- Use ‘formats’ instead of ‘generated’ [Sijis Aviles]

Other

- Refactor: Pass `gogoutils.Generator` to templates. [Sijis Aviles]

1.12.95 v3.50.2 (2018-07-24)

- Converting to None for default user. [Joel Vasallo]
- Support to run pipelines using a default *Run As User* [Joel Vasallo]
 - Needed for Fiat enabled applications to specify service account

1.12.96 v3.50.1 (2018-07-24)

- Test: Fix failed linting checks. [Sijis Aviles]

1.12.97 v3.50.0 (2018-07-24)

- Converting logic to jsonify list. [Joel Vasallo]
- Added missing keys to test. [Joel Vasallo]
- Added init support for Fiat Permissions and Traffic Guards. [Joel Vasallo]
- Added initial Python 3.7 support. [Joel Vasallo]
- Test: Validate EC2 Pipeline Type configuration. [Nate Tangsurat]
- Refactor: Support overriding EC2 Pipeline Types. [Nate Tangsurat]

Need a way to inject custom Pipeline Types that are deployed with EC2.

1.12.98 v3.49.0 (2018-06-19)

Fix

- Travis deploy Tags. [Nate Tangsurat]

Travis has a small quirk where one deploy block cannot specify both Branches and Tags. To get around this, there needs to be one block for the master Branch, and one block for Tags.

Features

- Deploy with Travis. [Nate Tangsurat]

1.12.99 v3.48.0 (2018-06-19)

Fix

- Lookup formats configuration in dict fashion. [Nate Tangsurat]
- Return dict form of configurations. [Nate Tangsurat]

ConfigParser can be accessed like a dictionary already, but converting the Type will help avoid more confusion.

- Avoid ConfigParser squashing. [Nate Tangsurat]

When adding values to a ConfigParser, everything gets squashed into a string. Avoid this by preserving the Python object as a whole.

Other

- Docs: Update configuration finder docstring. [Nate Tangsurat]
- Docs: Update docstring for formats extraction. [Nate Tangsurat]
- Refactor: Access configs like dictionary. [Nate Tangsurat]
- Docs: Fixed wrong key in lambda_environment example for application json. [MilosRasic]

1.12.100 v3.47.0 (2018-06-19)

- Various files modified to update copyright message to 2018. [Alex King]

The following script was used to accomplish this:

```
for file in grep -R "Copyright 2016" .|cut -f1 -d:|sort -u; do sed 's/Copyright 2016/Copyright 2018/g' < $file >$file.new; mv $file.new $file; done
```

1.12.101 v3.46.2 (2018-06-19)

Fix

- Not every Stage has a status. [Nate Tangsurat]
- Also skip over NOT_STARTED Stages. [Nate Tangsurat]
- SpinnakerTaskError skip over good Stages. [Nate Tangsurat]

Features

- DEBUG Stages in SpinnakerTaskError. [Nate Tangsurat]

Other

- Refactor: Avoid extra union in loop. [Nate Tangsurat]

1.12.102 v3.46.1 (2018-05-22)

- Added support for AWS Xray in Services. Added missing documented services. [Joel Vasallo]
- Added ImportError logging issues with reading dynamic config. [Joel Vasallo]
- Added Generic IAM JSON policy for foremast. [Joel Vasallo]
- Updated AWS Infra assumptions with basic IAM Policy. [Joel Vasallo]

1.12.103 v3.46.0 (2018-05-02)

- Removed hardcoded versions and updated spinnaker doc links. [Joel Vasallo]
- Fixed test case count with removal of default bucket name. [Joel Vasallo]
- Cleaned up example bucket name in templates. [Joel Vasallo]
- Fixed documentation syntax for image.rest. [Joel Vasallo]
- Updated docs and logging. [Joel Vasallo]
- Added docs for bake_instance_type. [Joel Vasallo]
- Added support for overriding aws_instance_type in Rosco via bake_instance_type. [Joel Vasallo]
- Added support for S3 bucket lifecycle configurations. [Joel Vasallo]

1.12.104 v3.45.1 (2018-04-16)

- Fixed lint error” [Doug Campbell]
- Fixed issue with non-east buckets already existing and giving errors. [Doug Campbell]

1.12.105 v3.45.0 (2018-04-16)

- Added line to fix lint errors. [Doug Campbell]
- Added Docs for new S3 enhancements. [Joel Vasallo]
- Fixed linting errors. [Joel Vasallo]
- Removed extra whitespace. [Joel Vasallo]
- Converted bucket policy to add or delete as well. [Joel Vasallo]
- Added support for bucket_encryption bucket_cors and revamped website. [Joel Vasallo]
 - Previously we would not clean up well if users toggle setting on or off
- Moved tagging, versioning and logging to master bucket settings. [Joel Vasallo]
 - In addition, exposed adding tags through app-json - Remove get tags?
 - Versioning toggles on and off based on enabled/disabled
 - Tagging toggles and adds/removes as you remove from json files
 - Logging toggles on and off based on enabled. - Default prefix? (<account>/<region>/<app>/)
- Changed logic on bucket versioning to be disabled by default. [Joel Vasallo]
 - Bucket versioning now correctly toggles versioning on or off based on settings in application json
- Fixed json error in configs. [Joel Vasallo]
- Added support for CORS on s3 bucket. [Joel Vasallo]
- Added logging_config to put_bucket_logging on buckets. [Joel Vasallo]
- Converted logging and versioning to similar structure. [Joel Vasallo]

- Added base configs for logging and versioning s3 buckets. [Joel Vasallo]

1.12.106 v3.44.0 (2018-04-16)

- Updated runner to pass primary region. [Doug Campbell]
- Updated s3deploy to use new bucket naming for deployments. [Doug Campbell]
- Fixed lint errors. [Doug Campbell]
- Added primary region and ability to generate buckets in other regions. [Doug Campbell]

1.12.107 v3.43.0 (2018-04-10)

- Added docs. [Doug Campbell]
 - Added unit tests. [Doug Campbell]
 - Fixed if logic and lint errors. [Doug Campbell]
 - Fixed logic on generating cli cmd. [Doug Campbell]
 - Changed flat to mirror. [Doug Campbell]
 - Fixed lint errors. [Doug Campbell]
 - Added cmd function, works for content header uploads too. [Doug Campbell]
 - Added logic for uploading s3 content in a flat structure instead of version. Does not work for metadata uploads yet. [Doug Campbell]
 - Updated Athena and Glue permissions. [Joel Vasallo]
 - Athena now includes glue:Get* as this is now required once Athena's data catalog is upgraded.
 - Glue was converted to generic Get* as this might need to be tweaked in the future
 - Docs: node8x, additional information. [Steven Basgall]
 - Docs: add all runtime strings. [Steven Basgall]
 - Updated scaling policy example. [Joel Vasallo]
 - Updated the docs to properly document period_minutes. [Joel Vasallo]
 - Added support for instance_warmup. [Joel Vasallo]
 - Delay between new instance coming up and metrics feeding into ASG
- Aggregate metrics

1.12.108 v3.42.0 (2018-04-02)

Fix

- Merge application and default security group rules. [Sijis Aviles]

This fixes a scenario where application.json is adding a rule definition that already exists in the ‘default_securitygroup_rules’ configuration. In those cases, the default settings would override what application.json defined.

We now combine both definitions into a single rule set.

Other

- Test: Validate merge of default security group rules. [Sijis Aviles]
- Style: isort. [Sijis Aviles]

1.12.109 v3.41.0 (2018-03-15)

Fix

- Use fully qualified path to file. [Sijis Aviles]

Other

- Style: Rename jinja variables. [Sijis Aviles]
- Refactor: Use pathlib for all template paths. [Sijis Aviles]
- Refactor: Use Pathlib objects until we need real path. [Sijis Aviles]

1.12.110 v3.40.2 (2018-02-28)

- Use default region in templates if not provided in configs. [Doug Campbell]
- Overrides regions in templates with regions in configs. [Doug Campbell]

1.12.111 v3.40.1 (2018-02-28)

- Fixed s3 app dns ttl issue with region specific configs. [Doug Campbell]

1.12.112 v3.40.0 (2018-02-26)

Fix

- Remove reference to invalid stage. [Sijis Aviles]

Other

- Fixed docs and typos. [Doug Campbell]
- Added region to s3, lambda, and datapipeline pipelines. [Doug Campbell]
- Added region specific docs. [Doug Campbell]
- Fixed pipeline tests. [Doug Campbell]
- Pipeline to use region specific config. [Doug Campbell]

1.12.113 v3.39.0 (2018-02-22)

- Updated s3 and sgs. [Doug Campbell]
- Fixed data pipeline unit tests. [Doug Campbell]
- Updated to use region specific configs. [Doug Campbell]

1.12.114 v3.38.1 (2018-02-23)

Fix

- Suppress error if file is not found. [Sijis Aviles]
This keeps existing behaviour with the previous library.

Other

- Docs: Better descriptive error. [Sijis Aviles]

1.12.115 v3.38.0 (2018-02-23)

Fix

- Hard check for commit id. [Sijis Aviles]
- Warn regarding gitlab lookup. [Sijis Aviles]
- Initialize project upon object creation. [Sijis Aviles]
- Remove need to b64decode object. [Sijis Aviles]
- Move debug logging to function. [Sijis Aviles]
- Use project name as user. [Sijis Aviles]
- Add additional information regarding asg policy actions. [Sijis Aviles]
- Word typo. [Sijis Aviles]
- Remove warning of example security groups. [Sijis Aviles]

Other

- Style: Ignore init docstrings. [Sijis Aviles]
- Refactor: Rename method so its more appropriate. [Sijis Aviles]
- Refactor: Use FileLookup for getting files in gitlab. [Sijis Aviles]
- Docs: Update docstrings. [Sijis Aviles]
- Test: Update docstring to be more descriptive. [Sijis Aviles]
- Test: Remove unused import. [Sijis Aviles]
- Test: Use decorator form for patch.object. [Sijis Aviles]
- Test: Rename variable to be more descriptive. [Sijis Aviles]

- Test: Use simpler way to mock constant. [Sijis Aviles]
- Refactor: Remove project_id references. [Sijis Aviles]
- Test: Add fallback test for lookup. [Sijis Aviles]
- Test: Add json url lookup tests. [Sijis Aviles]
- Test: Fix ami lookup tests. [Sijis Aviles]
- Refactor: Git api call for getting ami file. [Sijis Aviles]
- Test: Fix git library references. [Sijis Aviles]
- Refactor: Switch library to support gitlab v4 api. [Sijis Aviles]

1.12.116 v3.37.1 (2018-02-22)

Fix

- Secondary region updating global dns record. [Sijis Aviles]

We should only update the global dns record when it matches the primary region. This addresses a problem where the secondary, region specific, domain was updating the global dns record.

Other

- Docs: Fix wrong parameter in docstring. [Sijis Aviles]

1.12.117 v3.37.0 (2018-02-22)

- Updated unittests to work with new properties format. [Doug Campbell]
- Fixed outputs to overlay region specific configs. [Doug Campbell]
- Updated to pass region into get_properties. [Doug Campbell]
- Fixed unittests with region configs. [Doug Campbell]
- Updated autoscaling and lambda to look at region specific configs. [Doug Campbell]

1.12.118 v3.36.0 (2018-02-20)

Features

- Allow setting lambda concurrency limit. [Will Kelleher]

1.12.119 v3.35.0 (2018-02-15)

- Fixed missing doc string. [Doug Campbell]
- Added unit tests for region specific configs. [Doug Campbell]
- Fixed linting errors. [Doug Campbell]
- Combined shared code from process_git_configs and process_runway_config into one function. [Doug Campbell]

- Added ability to override configs with region specific settings. [Doug Campbell]
- Updated templates to work with nested region blocks. Keeps backwards compatibility. [Doug Campbell]

1.12.120 v3.34.6 (2018-02-13)

- Disabled trigger for other pipelines besides just ec2. [Doug Campbell]

1.12.121 v3.34.5 (2018-02-08)

- Pass all pipeline data to wrapper creation. [Doug Campbell]

1.12.122 v3.34.4 (2018-01-19)

- Test: Add pipeline runner tests. [Sijis Aviles]
- Refactor: Fallback to default pipeline for unknown types. [Sijis Aviles]
This allows users to create custom pipeline type and not need to make code updates for those new types to work.
- Docs: Fix header for types option. [Sijis Aviles]
- Refactor: Remove defined types in pipeline_wrapper. [Sijis Aviles]
We already have defined Pipeline types for lambda, s3 and datapipeline.

1.12.123 v3.34.3 (2018-01-11)

- Fixed issue where v2 keys may not be in tlscert template. [Joel Vasallo]

1.12.124 v3.34.2 (2018-01-11)

Fix

- We should default the env to current one. [Sijis Aviles]

1.12.125 v3.34.1 (2018-01-11)

Fix

- Remove default start/end port. [Sijis Aviles]

We should fail if the ports are not specified.

Other

- Style: lint fixes. [Sijis Aviles]
- Refactor: Split create_security_group to be more testable. [Sijis Aviles]

1.12.126 v3.34.0 (2018-01-10)

Fix

- Cross account SG within same account. [Sijis Aviles]

We need to ensure that we do not attempt cross account security lookups when the account is the same as the app. We should only trigger that lookup when the app and SG are different.

1.12.127 v3.33.0 (2018-01-10)

- Fixed docstrings and unneeded imports. [Joel Vasallo]
- Fixed test case for v1. Multi-region not supported in v1. [Joel Vasallo]
- Added test cases for testing both v1 and v2 template format for tls certs. [Joel Vasallo]
- Cleaned up code using yapf. [Joel Vasallo]
- Added new tlscert_naming format to allow support for ACM. [Joel Vasallo]
 - ACM creates certs per region, this change enables users to reference certs by friendly name instead of long ARN.

1.12.128 v3.32.0 (2018-01-10)

- Formatting fixes. [Joel Vasallo]
- Removed unneeded import. [Joel Vasallo]
- Added docs for athena and glue. [Joel Vasallo]
- Added Athena/Glue IAM permissions to base templates. [Joel Vasallo]

1.12.129 v3.31.0 (2017-12-20)

- Updated test cases to handle mock properly. [Joel Vasallo]
- Added sample test case. [Joel Vasallo]
- Moved security_group duplicate logic to a utility function. [Joel Vasallo]
- Fixed an issue with security group local variable. [Joel Vasallo]
- Fixed formatting to be compliant with standards. [Joel Vasallo]
- Added support for removing duplicate security group references based on name. [Joel Vasallo]
- Added pypi badge. [Doug Campbell]

1.12.130 v3.30.0 (2017-12-19)

- Fixed lint error. [Doug Campbell]
- Added ability to specify if scale-down policy should be added. [Doug Campbell]
- Fixed lint error. [Doug Campbell]
- Added docs and missing trust template for rolling pipeline. [Doug Campbell]

1.12.131 v3.29.0 (2017-12-14)

Fix

- Add logging on which CIDR we have checked. [Sijis Aviles]
 - Remove security group network restrictions. [Sijis Aviles]
- Closes #232

Other

- Added missing : [Steven Basgall]
- Additional ARN verification. [Steven Basgall]
- Return ARN directly if topic name appears to be an ARN. [Steven Basgall]
- Fixed linting error with consts. [Doug Campbell]
- Added rolling ec2 pipeline type and made it work same as ec2 (just different templates) [Doug Campbell]

1.12.132 v3.28.1 (2017-12-04)

Fix

- Return Task status. [Nate Tangsurat]
- Better external template error message. [Nate Tangsurat]
- Better ERROR when template not found. [Nate Tangsurat]

Other

- Removed hardcoded image size limit, this can be done in the templates. [Doug Campbell]
- Docs: Update Task docstrings. [Nate Tangsurat]
- Docs: Update get VPC ID docstring. [Nate Tangsurat]
- Docs: Specify exceptions raise by templates. [Nate Tangsurat]
- Docs: Simplify docstring. [Nate Tangsurat]

1.12.133 v3.28.0 (2017-09-19)

Fix

- Multiple lambda s3 trigger event creation. [Will Kelleher]

1.12.134 v3.27.0 (2017-09-14)

- Explicitly promote canary to latest. [Doug Campbell]
- Updated docs for alpha deployments. [Doug Campbell]
- Cleaned up promote logic. [Doug Campbell]
- Added alpha pathing for s3 deployments. [Doug Campbell]
- Fixed lint error. [Doug Campbell]
- Combined sync_to functions to remove duplicate code. [Doug Campbell]
- Moving S3 docs to proper section in application-master. [Joel Vasallo]
- Added note about friendly DNS record if website enabled. [Joel Vasallo]
- Added note about path and limitations around slashes. [Joel Vasallo]
- Removed confusing default references and fixed options for bucket_acl. [Joel Vasallo]
- Initial commit of S3 docs for OSS release. [Joel Vasallo]

1.12.135 v3.26.0 (2017-09-07)

- Changed references to metadata to make it more abstract. [Doug Campbell]
- Fixed lint errors. [Doug Campbell]
- Clean up and doc strings. [Doug Campbell]
- Reorganized into ability to specify path instead of searching. [Doug Campbell]
- Check for content encoding directories and started to build cmds. [Doug Campbell]

1.12.136 v3.25.2 (2017-09-06)

- Added sorted to make a copy of the list so global variable is not mutated. [Doug Campbell]

1.12.137 v3.25.1 (2017-08-30)

Fix

- Use better variable name. [Sijis Aviles]
- Ensure security is create prior to adding rules. [Sijis Aviles]

This should fix a dependency loop where default security group was defined in config.py as using a self reference. In those instances you could not create a security group because its self-reference did not exist.

Other

- Test: Fixup tests with recent refactoring. [Sijis Aviles]
- Refactor: Separate creating security group method. [Sijis Aviles]

1.12.138 v3.25.0 (2017-08-28)

Features

- \$self in security group config resolves to application. [Sijis Aviles]

Other

- Style: Fix import order. [Sijis Aviles]
- Test: Preserve rules during self-reference update. [Sijis Aviles]
- Refactor: Simplify logic for self references. [Sijis Aviles]
- Docs: Add note about \$self in security group rules. [Sijis Aviles]
- Docs: Use correct header styling. [Sijis Aviles]
- Test: Add security group self reference test. [Sijis Aviles]
- Test: Consolidate all security group tests. [Sijis Aviles]
- Added pipeline config instance links docs. [Doug Campbell]

1.12.139 v3.24.0 (2017-08-25)

Fix

- Remove reason of error after assert statement, fix documentation on unittests. [Varun Rao]
- Fix issues from code review. [Varun Rao]
- Add docstrings and fix naming error in consts. [Varun Rao]
- Change pipeline_configs to pipeline_config. [Varun Rao]
- Additions from pull_request. [Varun Rao]
- Iterate through values of Links instead of keys. [Varun Rao]
- Change pipeline_configs to pipeline_config. [Varun Rao]
- Additions from pull_request. [Varun Rao]
- Change name of links. [Varun Rao]
- Changed logic to accept blank filter pattern. [Steven Basgall]

Other

- Tests: Unit Tests for app creation. [Varun Rao]
- Feature: add instance links to pipeline configs. [Varun Rao]
- Docs: docstrings on unittests. [Varun Rao]
- Tests: Unit Tests for app creation. [Varun Rao]
- Style: Rename spinnaker_links due to redundancy. [Varun Rao]
- Feature: Split up retrieval of template. [Varun Rao]

- Feature: add instance links to pipeline configs. [Varun Rao]
- Feature: Retrieve links. [Varun Rao]
- Updated rds-db example. [Joel Vasallo]
removed confusing placeholder text in example usage
- Updated docs to reflect change from dict to list and updated template generation. [Joel Vasallo]
- Preventing setting of username in RDS-DB. [Joel Vasallo]
 - Forces App Name to be used for the user name to prevent folks from specifying any user.
- Docs: Clean up when changing ELB purpose. [Nate Tangsurat]

1.12.140 v3.23.1 (2017-08-22)

Fix

- Explicitly encode lambda trigger rules. [Will Kelleher]
- Add debugging to aws responses. [Sijis Aviles]
- Use correct variable name. [Sijis Aviles]
- Use proper constant variable name. [Sijis Aviles]
- Use better way to string check a string. [Sijis Aviles]
- Remove unused methods. [Sijis Aviles]
- Disable invalid-name pylint warning. [Sijis Aviles]
- Remove unused import. [Sijis Aviles]

Other

- Style: Supress warning on invalid name. [Sijis Aviles]
- Style: Remove lint supression. [Sijis Aviles]
- Style: Supress a couple of warnings. [Sijis Aviles]
- Style: yapf the file. [Sijis Aviles]
- Refactor: s/update_default_securitygroup_rules/update_default_rules. [Sijis Aviles]
- Refactor: s/construct_pipeline_block_datapipeline/construct_datapipeline. [Sijis Aviles]
- Refactor: s/configure_load_balancer_attributes/configure_attributes. [Sijis Aviles]
- Refactor: Simply logic and set single return line. [Sijis Aviles]
- Refactor: for/else conditions so it conforms style. [Sijis Aviles]
- Style: Use uppercase variable name. [Sijis Aviles]
- Style: Disable complex methods mccabe warnings. [Sijis Aviles]
- Chore: Silence several pylint warnings. [Sijis Aviles]
- Docs: Update docstrings. [Sijis Aviles]
- Chore: Remove runtests.py. [Sijis Aviles]

- Chore: Ensure linting failures fail ci. [Sijis Aviles]
- Docs: Reflow. [Nate Tangsurat]
- Docs: Rework workflow. [Nate Tangsurat]
- Docs: Reflow and consistency. [Nate Tangsurat]
- Docs: Reflow. [Nate Tangsurat]
- Docs: Reflow. [Nate Tangsurat]
- Docs: Fix section line lengths. [Nate Tangsurat]
- Docs: Fix INI syntax highlighting. [Nate Tangsurat]
- Docs: Reflow and fix section line lengths. [Nate Tangsurat]
- Docs: Fix section underline lengths. [Nate Tangsurat]
- Docs: Clean lambda environment JSON example. [Nate Tangsurat]
- Docs: Reflow and clean up JSON examples. [Nate Tangsurat]
- Docs: Use JSON data types. [Nate Tangsurat]
- Docs: Update S3 service example. [Nate Tangsurat]
- Docs: Fix S3 service wording. [Nate Tangsurat]
- Style: Reflow text. [Nate Tangsurat]

1.12.141 v3.23.0 (2017-08-11)

Fix

- Remove prod, prods, and prodp for public templates. [Varun Rao]
- Fix formatting in json. [Varun Rao]
- Include parameter name in function call. [Varun Rao]
- Remove prod, prods, and prodp for public templates. [Varun Rao]
- Add docstrings for pipeline_configs. [Varun Rao]
- Remove unneeded parenthesis around conditional. [Sijis Aviles]
- Simplify conditional check. [Sijis Aviles]
- Remove unused variables. [Sijis Aviles]
- Remove unused import. [Sijis Aviles]
- Use better variable name. [Sijis Aviles]
- Better variable names. [Nate Tangsurat]
- Raise proper error. [Nate Tangsurat]
- Lazy logging. [Nate Tangsurat]
- Passed in arguments are for future use. [Nate Tangsurat]

Features

- Include chaos_monkey configs. [Varun Rao]
- Include chaos_monkey configs. [Varun Rao]
- Include chaos_monkey configs. [Varun Rao]
- Change order of create_app and write_configs. [Varun Rao]
- Retrieve pipeline.json dict, pass pipeline.json dictionary to create_app. [Varun Rao]

Other

- Docs: Fix spelling in description of exceptions. [Varun Rao]
- Docs: Change min time to minimum_time. [Varun Rao]
- Docs: Add units to mean and min time. [Varun Rao]
- Docs: ADD explanation of Mean Time. [Varun Rao]
- Docs: Remove new line. [Varun Rao]
- Docs: Sphinx documentation for chaos monkey. [Varun Rao]
- Docs: Remove new line. [Varun Rao]
- Docs: Sphinx documentation for chaos monkey. [Varun Rao]
- Docs: Remove new line. [Varun Rao]
- Docs: Sphinx documentation for chaos monkey. [Varun Rao]
- Docs: Remove new line. [Varun Rao]
- Docs: Sphinx documentation for chaos monkey. [Varun Rao]
- Style: Fix formatting of json. [Varun Rao]
- Style: Properly give unused variables a more obvious name. [Sijis Aviles]
- Style: Fix variable names. [Sijis Aviles]
- Style: Fix variable names. [Sijis Aviles]
- Style: Fixup comments. [Sijis Aviles]
- Style: Use lazy logging. [Sijis Aviles]
- Style: Fixup comments. [Sijis Aviles]
- Style: Help with linting. [Sijis Aviles]
- Docs: Update docstrings. [Nate Tangsurat]
- Docs: Update S3 deployment docstrings. [Nate Tangsurat]
- Style: Pylint configuration. [Nate Tangsurat]
- Test: Ignore docs. [Nate Tangsurat]
- Style: First run through with isort. [Sijis Aviles]
- Chore: Add isort config. [Sijis Aviles]
- Style: First run through with yapf. [Sijis Aviles]
- Chore: Add yapf config. [Sijis Aviles]

1.12.142 v3.22.5 (2017-08-07)

Fix

- Remove duplicate variable definition. [Sijis Aviles]
- Ensure we handle dynamic configs properly for default_securitygroup_rules. [Sijis Aviles]
- Move config entry to [base] [Steven Basgall]

Features

- Added example configs. [Steven Basgall]

Other

- Docs: Update key to match other examples. [Sijis Aviles]
- Docs: Add docstring. [Sijis Aviles]
- Tests: Fix missing mock and rename variable. [Sijis Aviles]
- Tests: Default security group tests. [Sijis Aviles]
- Docs: Fix example for default_securitygroup_rules. [Sijis Aviles]
- Test: added test and slight refactor for test. [Steven Basgall]
- Added const tests. [Steven Basgall]
- Doc: default security group examples. [Steven Basgall]
- Style: indent whitespace. [Steven Basgall]
- Added ability to have default application securitygroup rules. [Steven Basgall]

1.12.143 v3.22.4 (2017-08-03)

- Add unit test and refactor. [Bekzot Azimov]
- Fix tests and refactor the code. [Bekzot Azimov]
- Add tagging to security groups. [Bekzot Azimov]

1.12.144 v3.22.3 (2017-07-31)

Fix

- Should not assign from an append operation. [Nate Tangsurat]

Other

- Test: Split IAM template tests with parametrize. [Nate Tangsurat]

See also: #208

- Test: Use generator for IAM template names. [Nate Tangsurat]
See also: #208
- Docs: Fix return description for Policy render. [Nate Tangsurat]
- Tests: IAM Policy templates should render. [Nate Tangsurat]
- Refactor: Extract local templates path. [Nate Tangsurat]
- Refactor: Extract IAM policy template rendering. [Nate Tangsurat]

1.12.145 v3.22.2 (2017-07-31)

Fix

- Create S3 lambda trigger SID with hash of rules. [Will Kelleher]

1.12.146 v3.22.1 (2017-07-28)

Fix

- Allow access to DynamoDB table/name/stream/ [Nate Tangsurat]

1.12.147 v3.22.0 (2017-07-28)

- Changed to sorted. [Doug Campbell]
- Create new list to avoid mutation errors. [Doug Campbell]
- Added in pipeline passing. [Doug Campbell]
- Unit test. [Doug Campbell]
- Removed kwarg. [Doug Campbell]
- Removed pipeline_type stuff from create_pipeline. not needed. [Doug Campbell]
- Fixed some typos and restructured get_template_name. [Doug Campbell]
- Changed kwargs to named arguments for clarity. [Doug Campbell]
- Flattened template block and moved formatting to one call. [Doug Campbell]
- Ran yapf. [Doug Campbell]
- Refactored construct blocks and create pipeline to segregate ec2 specific logic. [Doug Campbell]
- Style: YAPF. [Nate Tangsurat]
- Style: YAPF. [Nate Tangsurat]

1.12.148 v3.21.2 (2017-07-26)

Fix

- Invalid trailing comma. [Nate Tangsurat]

1.12.149 v3.21.1 (2017-07-25)

- Pass all pipeline data to templates. [Doug Campbell]

1.12.150 v3.21.0 (2017-07-24)

Fix

- Add check to existing conditional. [Sijis Aviles]
- Typos and docstring example. [Sijis Aviles]
- Remove correct item from list. [Sijis Aviles]
- Detect value as native python type. [Sijis Aviles]
- Remove empty items from security group list. [Sijis Aviles]

Features

- Make security groups for elb/ec2 environment specific. [Sijis Aviles]

Other

- Refactor: Use set() earlier in iteration. [Sijis Aviles]
- Style: Update docstring. [Sijis Aviles]
- Refactor: Use different way to strip out empty entries. [Sijis Aviles]
- Style: Word wrap to 80 chars. [Sijis Aviles]
- Docs: Add json example for security groups. [Sijis Aviles]
- Style: Quick lint fix. [Sijis Aviles]
- Tests: Validate passing dictionary from a dynamic config. [Sijis Aviles]
- Refactor: Populate security groups for each environment. [Sijis Aviles]
- Test: Fix elb test case to include an environment. [Sijis Aviles]
Our default environment is just an empty string and this is a workaround for that problem.
- Refactor: Use new security group format. [Sijis Aviles]
- Style: Linting fixes. [Sijis Aviles]
- Test: Add tests for security group changes. [Sijis Aviles]

1.12.151 v3.20.5 (2017-07-20)

- Simplify the logic of updating tags and add tests. [Bekzot Azimov]
- Separate out get tags logic. [Bekzot Azimov]
- Keep old s3 tags when new tags added. [Bekzot Azimov]

1.12.152 v3.20.4 (2017-07-11)

- Add limited cloudwatch iam access services option. [Bekzot Azimov]
- Add limited cloudwatch iam access services option. [Bekzot Azimov]

1.12.153 v3.20.2 (2017-07-11)

Fix

- Include seconds for timeout failure. [Nate Tangsurat]
- Use custom exception for inconclusive Task. [Nate Tangsurat]
- Remove exception `pass` statements. [Nate Tangsurat]
- Fail when Project ID lookup fails. [Nate Tangsurat]
- Bucket name, logging, access group as property. [Steven Basgall]

Features

- New Check Task inconclusive exception. [Nate Tangsurat]
- Exception for GitLab API errors. [Nate Tangsurat]

Other

- Fixed bug with pipelineId in data pipeline. [Doug Campbell]
- Docs: Fix test docstring. [Nate Tangsurat]
- Tests: Refactor duplicate test. [Nate Tangsurat]
- Tests: Remove unused mock. [Nate Tangsurat]
- Tests: Asserts outside the exception context. [Nate Tangsurat]
- Tests: Don't wait so long to fail. [Nate Tangsurat]
- Tests: Move retry task to new module. [Nate Tangsurat]
- Tests: Update to new Task inconclusive exception. [Nate Tangsurat]
- Tests: Update to new inconclusive Task exception. [Nate Tangsurat]
- Tests: Faster Check Task with new keyword argument. [Nate Tangsurat]
- Refactor: Move Check Task wait to argument. [Nate Tangsurat]
- Tests: Inconclusive Check Task polling. [Nate Tangsurat]
- Tests: Check Task keep polling. [Nate Tangsurat]
- Tests: Check Task failure. [Nate Tangsurat]
- Tests: Check Task success functionality. [Nate Tangsurat]
- Docs: Add periods to docstrings. [Nate Tangsurat]
- Tests: Check for GitLab API error during lookup. [Nate Tangsurat]
- Read group a different way. [Steven Basgall]

- Initial commit of s3 tagging. [Steven Basgall]

1.12.154 v3.20.1 (2017-06-29)

- Fixed variable name. [Doug Campbell]
- Added logic for handling pipeline ID of onetime pipelines. [Doug Campbell]
- Added runway_dir to fix errors in onetime pipelines. [Doug Campbell]

1.12.155 v3.20.0 (2017-06-28)

Fix

- Use dirty tag when developing. [Nate Tangsurat]
- Remove local scheme from version. [Nate Tangsurat]
Required by PEP 440 when uploading to PyPI.
- Inconsistent quoting in PR. [Steven Basgall]
- Skip creation of resource id if resource id already exists. [Steven Basgall]
- Flip statement IDs to reflect http method permission. [Steven Basgall]
- Ensure no in SID. [Steven Basgall]
- Don't eMail on Travis success. [Nate Tangsurat]

Features

- On function update also update tags. [Steven Basgall]

Other

- Added canary deploy_type to the docs. [Doug Campbell]
- Refactored path creation. [Doug Campbell]
- Removed canary specific deployment stage for base templates. [Doug Campbell]
- Better logging. [Doug Campbell]
- Updated templates for s3 deployments. [Doug Campbell]
- Ran yapf on s3deploy. [Doug Campbell]
- Added sync_to_canary functionality. [Doug Campbell]
- Remove spaces around non-assignment = [Steven Basgall]
- Refactor: removed redundant call to attach_method. [Steven Basgall]
- Refactor: add more specific names to arn variables and remove duplicate variable declarations. [Steven Basgall]
- Added passing of source ARN to permissions in lambda function policy and add function policies to alias and unqualified functions. [Steven Basgall]
- Docs: ASG Health Check type override note. [Nate Tangsurat]

- Docs: Note on ASG Health Check in dev environment. [Nate Tangsurat]
- Style: Strip. [Nate Tangsurat]
- Docs: Fix typo. [Sijis Aviles]
- Docs: Fix extra character. [Nate Tangsurat]
- Docs: Fix Scaling Policy key name. [Nate Tangsurat]
- Docs: Fix wrong key in pipeline.json. [Sijis Aviles]

1.12.156 v3.19.3 (2017-06-09)

- MR feedback, removed else. [Doug Campbell]
- Added paginating for finding lambda function arn. [Doug Campbell]
- Updated data pipeline docs. [Doug Campbell]
- Added ability to activate pipeline on deployment. [Doug Campbell]
- Test: Simplify travis config. [Sijis Aviles]

1.12.157 v3.19.2 (2017-06-02)

- Move tagging to create lambda from update. [Bekzot Azimov]

1.12.158 v3.19.1 (2017-05-31)

- Updated docs and renamed template for rds-db. [Joel Vasallo]
- Added docs for rds-db and updated refs to resource-id. [Joel Vasallo]
- Added support for rds-db iam policy. [Joel Vasallo]
- Fixed typo in elasticache iam policy. [Joel Vasallo]
- Update index.rst. [Joel Vasallo]
- Added data pipeline config docs. [Doug Campbell]

1.12.159 v3.19.0 (2017-05-31)

- Fixed function doc return types. [Doug Campbell]
- Removed data pipeline main. We do not need two entry points. [Doug Campbell]
- Continued working on bad exceptions. [Doug Campbell]
- Added return true for easier unit testing. [Doug Campbell]
- Added better error handling around bad json. [Doug Campbell]
- Added exception for issues creating data pipeline. [Doug Campbell]
- Added assert to test case. [Doug Campbell]
- Fixed unittest with mocking paginate. [Doug Campbell]
- Feedback from #173. [Doug Campbell]

- Removed endif from jinja template. [Doug Campbell]
- Refactored unit tests. [Doug Campbell]
- Added ability to paginate through listing Data Pipelines to get all results. [Doug Campbell]
- Added unit tests. [Doug Campbell]
- Added datapipeline type and templates. [Doug Campbell]
- Added datapipeline functionality into runner. [Doug Campbell]
- Added tags to pipeline creation. [Doug Campbell]
- Added functionality to create pipeline and put definition. [Doug Campbell]
- Added datapipeline dir and defaults to config template. [Doug Campbell]

1.12.160 v3.18.0 (2017-05-31)

Fix

- Travis lint with matrix exclusions. [Nate Tangsurat]
Go back to four Build Jobs and run linting with Python 3.6.
- Run separate lint. [Nate Tangsurat]
- Run Travis without multiple environments. [Nate Tangsurat]

Features

- Show package version with `-version` [Nate Tangsurat]
- Generate changelog. [Nate Tangsurat]

Other

- Refactor: Unused import. [Nate Tangsurat]
- Refactor: Sphinx should also use `setuptools_scm` [Nate Tangsurat]
- Refactor: Use `setuptools_scm` for package version. [Nate Tangsurat]
- Docs: Forgot changelog. [Nate Tangsurat]
- Tests: Travis should use more Python versions. [Nate Tangsurat]
- Docs: Fix include path to configuration example. [Nate Tangsurat]
- Docs: Remove autogenerated docs. [Nate Tangsurat]

1.12.161 v3.17.4 (2017-05-31)

- Tag option for lambda deployments. [Bekzot Azimov]
- Docs: Add reference for intersphinx. [Nate Tangsurat]
- Added global prefix and continue loop instead of return. [Doug Campbell]
- Moved prefix into creating permission logic. [Doug Campbell]

1.12.162 v3.17.3 (2017-05-31)

Fix

- Indentation warnings. [Sijis Aviles]

Other

- Added support for alias policies. [Doug Campbell]
- Look for old Lambda policies too. [Doug Campbell]
- Added foremast-* to all SIDs and added functionality for cleaning up foremast created Lambda policies. [Doug Campbell]
- Docs: Fix ‘Title underline too short.’ warning. [Sijis Aviles]
- Added conditional to autoscaling templates to fix issue with decrease policies. [Doug Campbell]

1.12.163 v3.17.2 (2017-05-19)

Fix

- Make readthedocs work again. [Sijis Aviles]
- Add variable assignment to replace statement. [Steven Basgall]
- Change print to debug log. [Chris Kamradt]

Features

- Configurable task timeouts. [Chris Kamradt]
Added support for configuring timeouts per task per environment.

Other

- Added floor for scaling policy math to keep it all integers. [Doug Campbell]
- Test: add timeout testing for tasks. [Chris Kamradt]
- Refactor: rename timeout vars and add docs. [Chris Kamradt]
- Test: Update tests with new task methods. [Chris Kamradt]
- Refactor: enable task timeout config. [Chris Kamradt]

1.12.164 v3.17.1 (2017-05-11)

Fix

- Re-enable linting. [Sijis Aviles]

Other

- Fixed tests base on correct structure. [Doug Campbell]
- Docs: Add test docstring information. [Nate Tangsurat]
- Tests: Renuminate properly. [Nate Tangsurat]
- Fixed bug where miltiple branch stages did not renumerate properly. [Doug Campbell]
- Added default for canary. [Doug Campbell]

1.12.165 v3.17.0 (2017-05-10)

- Refactor: consolidate canary conditionals. [Steven Basgall]
- Initial concept for canary userdata in foremast. [Steven Basgall]

1.12.166 v3.16.1 (2017-05-04)

- Print out AWS exception so we can debug why this happens. [Doug Campbell]
- Place log statements variable references after variables. [Bekzot Azimov]

1.12.167 v3.16 (2017-05-02)

Fix

- Handle missing dynamic config file. [Sijis Aviles]
- Version number parser broke docs compilation. [Nate Tangsurat]
Version number uses *setuptools* which moved the *packaging* module up a level.

Features

- Add ability to use a dynamic config. [Sijis Aviles]
We are now able to load a python file ‘config.py’ as a source of configuration details.

Other

- Added config.py to gitignore. [Doug Campbell]
- Docs: Add details regarding config.py. [Sijis Aviles]
- Refactor: Loading static or dynamic configurations. [Sijis Aviles]

This ensures that a static config is preferred over a dynamic one. When a dynamic or static config is not found, use the default and show where we are searching for configs.

- Updated docs to have example of Security Group usage. [Joel Vasallo]

1.12.168 v3.15.1 (2017-04-21)

- Fixed invalid app cookie stickiness policy names. [Joel Vasallo]
 - Some JS libraries ship with dots by default in their cookie name, this at least strips it out.

1.12.169 v3.15.0 (2017-04-20)

- Fixed remaining bugs, working version with s3 creation and deployment. [Doug Campbell]
- Fixed issues with shared bucket naming, works with updated gogo-utils. [Doug Campbell]
- Fixed typos. [Doug Campbell]
- Added logic for making sure shared bucket exists. [Doug Campbell]
- Added logic for getting shared bucket name. [Doug Campbell]

1.12.170 v3.14.1 (2017-04-19)

- Fixed missing test case for create_elb. [Joel Vasallo]

1.12.171 v3.14.0 (2017-04-19)

- Added proper tag structure to RELEASING docs. [Joel Vasallo]
- Cleaned up logging statements and if conditional on modify-lb- attributes. [Joel Vasallo]
- Added docs for enabling access_log, connection_draining_timeout, and idle_timeout. [Joel Vasallo]
 - Also fixed logic to provide defaults in case of enabling or disabling of features
- Added support for ELB Access Logs. [Joel Vasallo]
- WIP: Enable modification of ELB attributes such as draining and timeouts. [Joel Vasallo]
- Adding recursive cp before s3 sync. [Bekzot Azimov]
- Added TODO to migrate ssh keyname generation out. [Joel Vasallo]
- Added docs to show migration of app_ssh_key to ssh_keypair. [Joel Vasallo]
- Added support to specify ssh_keypair to use. [Joel Vasallo]
- Removed linting from travis too. [Doug Campbell]
- Commented out linting to fix CI jobs until the PR in prospector is accepted. [Doug Campbell]
- Docs: Add python3.6 runtime option to lambda documentation. [Steven Basgall]

1.12.172 v3.12.2 (2017-04-14)

Fix

- Remove extra whitespace around config locations. [Sijis Aviles]
- Set section warning to info level. [Sijis Aviles]
- Ensure proper logging format when validating config. [Sijis Aviles]

- Remove duplicate key. [Sijis Aviles]

Features

- Add SimpleDB IAM template. [Will Kelleher]
- Add multipart upload permissions to s3 template. [Will Kelleher]

Other

- Fix 120 line character limit. [Doug Campbell]
- Added ‘–exact-timestamps’ to s3 sync. [Doug Campbell]
- Increased visibility how grace_period is calculated. [Joel Vasallo]
- Added support for app_grace_period to extend grace_period. [Joel]
- Added support for modifying ASG grace period from json files. [Joel]
- Updated docs to reflect new services added to foremast. [Joel]
- Added support for kinesis in json. [Joel]
- Ensure threshold and period_minutes are integers. [Sijis Aviles]

This handles cases where those keys are set as strings even though we expect integers.

Fixes #100

- Chore: FIXME to handle pyapi-gitlab return False. [Nate Tangsurat]

1.12.173 v3.12.1 (2017-03-13)

- Fixed formatting. [Doug Campbell]
- Catches problems with empty or missing artifact path. [Doug Campbell]

1.12.174 v3.11.0 (2017-03-10)

Fix

- Removed unneeded subnet config. [Steven Basgall]
- Missing close paren. [Steven Basgall]

Features

- Template additions for judgement promote. [Steven Basgall]
- New pipeline stage ordering. [Steven Basgall]
- Added artifact version to deploy-s3 stage template. [Steven Basgall]
- Changed default jenkins job name for s3 deploy stage. [Steven Basgall]
- Code changes to utilize new pipeline templates. [Steven Basgall]
- Added s3 spinnaker templates. [Steven Basgall]

Other

- Changed if to elif in jinja2 template. [Doug Campbell]
- Removed unneeded import. [Steven Basgall]
- Correct missing S3 strings. [Steven Basgall]
- Added type for s3 in consts.py. [Steven Basgall]
- Added s3 pipeline block - repeatable stage for s3. [Steven Basgall]
- Added s3 conditional to pipeline wrapper. [Steven Basgall]
- Added configs to runner s3 jobs. [Doug Campbell]
- Added bucket DNS. [Doug Campbell]
- Changed bucket policy to empty dict. [Doug Campbell]
- Added S3 website settings for meeting a bucket serve static content. [Doug Campbell]
- Removed lingering print statement. [Doug Campbell]
- Implemented feedback from #136. [Doug Campbell]
- Renabled create_bucket. [Doug Campbell]
- Added s3 deployments to runner. [Doug Campbell]
- Switched for SH to subprocess to removed a dependency. [Doug Campbell]
- Changed s3 deployments to us SH and aws cli. [Doug Campbell]
- Added function to copy deployed version to LATEST. [Doug Campbell]
- Added function to upload all files and directories of artifact path. [Doug Campbell]
- Setup skeleton code for s3 deployments. [Doug Campbell]
- Fixed doc strings to better reflect what the class does” [Doug Campbell]
- Fixed formatting. [Doug Campbell]
- Added to runner, reorganized under s3 dir with archaius. [Doug Campbell]
- Added policy attachment and updated default configs to include S3. [Doug Campbell]
- Started skeleton for s3deployment infrastructure. [Doug Campbell]
- Updated policies block in elb templates. [Joel Vasallo]
- Cleaned up references to backendPolicies and fixed hardcoded test values. [Joel Vasallo]
- Updated docs, tests, and cleaned up if logic. [Joel Vasallo]
- Fixed test cases for missing mocks and added default response if dict key missing for backend policies. [Joel Vasallo]
- Defaulting to empty lists instead of NoneType elb_settings for policies. [Joel Vasallo]
- Added new elb_json structure for backendPolicies to test cases. [Joel Vasallo]
- Added test for testing backend policy addition. [Joel Vasallo]
- Updated log statements and docs. [Joel Vasallo]
- Updated references to null from empty list and updated base reference template. [Joel Vasallo]
- Formatted variable naming to be easier to read and cleaner. Added docs like a good developer. [Joel Vasallo]

- Added support for backend server policies in foremast config. [Joel Vasallo]
 - Needed for legacy support for WebSockets where they need to pass to backend systems through ProxyProtocol
 - Cleaned up naming of ‘policies’; to be specific to ‘listenerpolicies’ in the foremast config, but in order to not break existing functionality, merged policies + listenerpolicies
- Switch static names to args. [Bekzot Azimov]
- Username and emoji as an argument. [Bekzot Azimov]
- Give a name and icon for slackbot. [Bekzot Azimov]

1.12.175 v3.10.3 (2017-02-22)

Features

- Added null default value for app:lambda_environment. [Steven Basgall]

Other

- Added missing client for route53 delete func. [Bekzot Azimov]
- Style: Extended title underlines. [Steven Basgall]
- Docs: Added lambda_environment documentation and example. [Steven Basgall]
- Style: removed parens from import. [Steven Basgall]
- Style: use project column width. [Steven Basgall]
- Added test. [Steven Basgall]
- Renamed clobbering key name - deviating from API naming. [Steven Basgall]
- Added environment dict. [Steven Basgall]

1.12.176 v3.10.2 (2017-02-21)

- Fixed issue with cookie expiration period default. [Doug Campbell]

1.12.177 v3.10.1 (2017-02-17)

- Moved string format outside conditionals. [Doug Campbell]
- Added cookie_ttl support and fixed bug with duplicate stickiness policies. [Doug Campbell]

1.12.178 v3.10.0 (2017-02-16)

- Added back in code formating to docs for stickiness. [Doug Campbell]
- Updated docs to include elb cookie stickiness. [Doug Campbell]
- Added loadbalancer cookie stickiness support. [Doug Campbell]
- Added setup.cfg for bdist universal. [Bekzot Azimov]

1.12.179 v3.9.3 (2017-02-09)

- Fixed issue with poorly named keywords. [Doug Campbell]

1.12.180 v3.9.2 (2017-02-06)

- Adding asg_whitelist data. [Bekzot Azimov]
- Removing hard coded asg instance defaults. [Bekzot Azimov]

1.12.181 v3.9.1 (2017-01-30)

- Moved test until utils with other dns tests, split tuple into two values. [Doug Campbell]
- Made the dns test more explicit. [Doug Campbell]
- Removed duplicate test. [Doug Campbell]
- Added unit test for getting zone_ids. [Doug Campbell]
- Added test case for finding existing record. [Doug Campbell]
- Refactored into a generic find_existing_record function. [Doug Campbell]
- Fixed item variable name to be more consistent, changed to record. [Doug Campbell]
- Added a lookup to make sure that the primary record exists before creating the secondary failover record. I also updated the CNAME lookup code to use paginators to solve a future issue. [Doug Campbell]
- Removed comments and fixed templates mocking. [Doug Campbell]
- Mocked config values so that it works with local foremast configs. [Doug Campbell]

1.12.182 v3.9.0 (2017-01-17)

- Removed lingering line that I forgot about. [Doug Campbell]
- Updated docs and ran yapf on elb code. [Doug Campbell]
- Added cookie stickiness logic. [Doug Campbell]

1.12.183 v3.8.1 (2017-01-16)

- Added rstring and fixed bad variable name. [Doug Campbell]
- Added feedback from #114. [Doug Campbell]
- Added logic for finding and deleting old CNAME records when updating to multiregion. [Doug Campbell]

1.12.184 v3.8.0 (2017-01-06)

Fix

- Add rule name to INFO. [Nate Tangsurat]

Other

- Added emr policy. [Doug Campbell]
- Docs: Fix argument type hint. [Nate Tangsurat]
- Style: YAPF. [Nate Tangsurat]

1.12.185 v3.7.2 (2017-01-04)

Fix

- Fatal message text. [Sijis Aviles]

Other

- Add description to debug output. [Sijis Aviles]
- Add additional debug logging. [Sijis Aviles]
- Use non-alias arn for adding permissions. [Sijis Aviles]
- Docs: Fix docstring. [Nate Tangsurat]

See also: PSOBAT-2925

1.12.186 v3.7.1 (2017-01-03)

- Chore: Debug Role name message. [Nate Tangsurat]

1.12.187 v3.7.0 (2017-01-03)

Fix

- Provide default Lambda Role configuration. [Nate Tangsurat]

Features

- New Lambda IAM Role name generator. [Nate Tangsurat]
gogo-utils 1.5 provides a new IAM Role name generator for Lambda Functions.
See also: gogoair/gogo-utils#4
- Support custom Lambda Role. [Nate Tangsurat]

Other

- Test: Use custom Lambda Role when provided. [Nate Tangsurat]
- Docs: Note on custom Lambda Role. [Nate Tangsurat]

1.12.188 v3.6.1 (2016-12-29)

Fix

- Remove hardcoded IAM name fields. [Nate Tangsurat]
See also: gogoair/gogo-utils#4
- Requirements for docs. [Nate Tangsurat]
- Update log messages for Lambda configuration. [Nate Tangsurat]
- Better Pipeline deletion. [Nate Tangsurat]

Handle bad names from manual Pipelines and exit when Pipeline name is most likely bad.

Features

- New Pipeline delete exception. [Nate Tangsurat]

Other

- Docs: Fix errant character. [Nate Tangsurat]
- Docs: Fix underline lengths. [Nate Tangsurat]
- Docs: Populate release and version fields. [Nate Tangsurat]
- Add application name in successful pipeline message. [Sijis Aviles]

1.12.189 v3.6.0 (2016-12-19)

Fix

- Remove need to delete Pipeline when updating. [Nate Tangsurat]
When “id” is *null*, a new Pipeline is automatically created. When “id” is an actual ID, the Pipeline is updated with the new configuration.
- Delete Pipeline using new Pipeline values. [Nate Tangsurat]
- Allow Manual Pipeline application and name. [Nate Tangsurat]
Only use default generated Spinnaker application and Pipeline name when not defined in JSON.
- Use safe name for manual Pipeline. [Nate Tangsurat]

Features

- Normalize Pipeline name function. [Nate Tangsurat]
Spinnaker does not allow special characters in the Pipeline name.

Other

- Test: Look for Pipeline in an Application. [Nate Tangsurat]
- Test: Check Pipeline name normalizer. [Nate Tangsurat]

1.12.190 v3.5.1 (2016-12-13)

- Made it so that when DNS is created the global record and region specific is created. This is configurable but true by default. [Doug Campbell]

1.12.191 v3.5.0 (2016-12-12)

Fix

- Change subcommand from tester -> validator. [Sijis Aviles]
- Testers take in arguments for later use. [Nate Tangsurat]
- Rename to use “testers” [Nate Tangsurat]
- Move all tester runner to subcommand. [Nate Tangsurat]

Features

- Add Gate tester subcommand. [Nate Tangsurat]
- Test Gate connection. [Nate Tangsurat]
- Stub out tester command. [Nate Tangsurat]

Provide some manner of testing Spinnaker setup.

Other

- Keep names consistent using validate. [Sijis Aviles]
- Style: YAPF. [Nate Tangsurat]

1.12.192 v3.4.0 (2016-12-07)

- Split off into separate functions in the dns utils. This cleans up the code. [Doug Campbell]
- Removing loops for regions. [Doug Campbell]
- Fixed issue with selecting failover status in DNS. [Doug Campbell]
- Added the ability to configure default primary region and if to create failover DNS. [Doug Campbell]
- Updated runner to run multiregion setup if two are specified. [Doug Campbell]
- Ran yapf and isort. [Doug Campbell]
- Added elb_zone_id getting, template works. Working commit just needs cleanup. [Doug Campbell]
- Added region specific DNS and failover creation and template. Still WIP. [Doug Campbell]
- Added start for failover record creation. Needs fleshed out. [Doug Campbell]

- Added regions to gogoutils generator. [Doug Campbell]

1.12.193 v3.3.1 (2016-12-05)

Fix

- Include api details in docs. [Sijis Aviles]

Other

- Added purpose to get_subnet for elbs. [Doug Campbell]

1.12.194 v3.3 (2016-11-30)

Fix

- Update to use CliArgs namedtuple. [Nate Tangsurat]
- Use namedtuple for passing parsed args. [Nate Tangsurat]
- Use better name for unknown arguments. [Nate Tangsurat]
- Allow unknown arguments to pass through. [Nate Tangsurat]
- Support rebuild subcommand arguments. [Nate Tangsurat]
- Rebuild should fail before getting all apps. [Nate Tangsurat]
- Rebuild docstring formatting. [Nate Tangsurat]
- Onetime subcommand –all should be flag. [Nate Tangsurat]
- Point autoscaling to entry point. [Nate Tangsurat]
- Point rebuild to entry point. [Nate Tangsurat]
- Point Pipeline onetime to entry point. [Nate Tangsurat]
- Point pipeline app subcommand to entry point. [Nate Tangsurat]
- Pipeline subcommand title. [Nate Tangsurat]
- Allow no argument passing fallback. [Nate Tangsurat]
- Point infra subcommand to function. [Nate Tangsurat]
- Environment argument overrides \$ENV. [Nate Tangsurat]
- Show argument defaults. [Nate Tangsurat]
- Support ENV environment variable. [Nate Tangsurat]

Features

- Short logging format. [Nate Tangsurat]
- Project argument for rebuild subcommand. [Nate Tangsurat]
- Add autoscaling subcommand. [Nate Tangsurat]

- Add rebuild subcommand. [Nate Tangsurat]
- Add pipeline subcommand. [Nate Tangsurat]
- Add infra subcommand. [Nate Tangsurat]
- Start unified ‘foremast’ CLI. [Nate Tangsurat]

Other

- Merge pull request #86 from gogoair/feat/commands. [Sijis Aviles]
feat: Unified CLI commands
- Refactor: Move rebuild ALL check. [Nate Tangsurat]

1.12.195 v3.2.0 (2016-11-15)

- Merge pull request #102 from gogoair/r53-multiplezonehandler. [fadi- almasri]
Route53 DNS Enhancements
- R53 DNS Fixes. [Joel Vasallo]
 - Enabled support for multiple Hosted Zones (skipping over invalid private DNS zones)
 - Added try/except to potential boto3 call for creating R53 records
 - Enhanced logging around creating R53 records
- WIP: Attempting to fix multiple DNS zones error. [Joel Vasallo]
- Added docs to highlight firehose support. [Joel Vasallo]

1.12.196 v3.1.7 (2016-11-07)

- Fixed typo in firehose policy template. [Joel Vasallo]
- Updated tests to use the tests/ directory specifically. [Joel Vasallo]

1.12.197 v3.1.6 (2016-11-01)

Fix

- Missing square bracket. [Nate Tangsurat]

1.12.198 v3.1.5 (2016-11-01)

Fix

- New exception, ForemastConfigurationFileError. [Nate Tangsurat]
See also: #96
- Handle missing environment configuration. [Nate Tangsurat]

Other

- Updated docs to show ElasticSearch support. [Joel Vasallo]
- Updated ElasticSearch permissions. [Joel Vasallo]
 - Now supports multiple domains (list names)
 - Fixed issue for hidden undocumented IAM Permissions (such as es:ESHttpGet)
- Test: Raise missing SG environment configuration. [Nate Tangsurat]

1.12.199 v3.1.4 (2016-10-24)

- Updated tests to expect Discovery. [Doug Campbell]
- Changed default provider to discovery. [Doug Campbell]

1.12.200 v3.1.3 (2016-10-24)

Fix

- Delete manual Pipeline before creating. [Nate Tangsurat]
See also: #72
- Use filename for Pipeline name. [Nate Tangsurat]
See also: #72
- Use new class name FileLookup. [Nate Tangsurat]
See also: #72
- Path expansions need empty string default. [Nate Tangsurat]
See also: #72
- Pass runway directory to Pipeline. [Nate Tangsurat]
See also: #72
- Pass runway directory to lookup. [Nate Tangsurat]
See also: #72
- Add manual Spinnaker Pipeline Class to init. [Nate Tangsurat]
See also: #72

Features

- Save runway directory for Pipeline usage. [Nate Tangsurat]
See also: #72
- Add manual Spinnaker Pipeline type. [Nate Tangsurat]
See also: #72
- New manual Pipeline class. [Nate Tangsurat]
See also: #72

- Default Pipeline file list. [Nate Tangsurat]

See also: #72

Other

- Refactor: Extract Pipeline deletion. [Nate Tangsurat]

See also: #72

- Style: isort, YAPF. [Nate Tangsurat]

See also: #72

- Docs: Add manual Pipeline settings. [Nate Tangsurat]

See also: #72

- Changed pipeline_id to be None if no existing pipeline exists. The empty string causes problems in S3. [Doug Campbell]

- Changed to pypi references. [Doug Campbell]

1.12.201 v3.1.2 (2016-10-20)

- Added logic to make sure the policies actually exist in ASG before trying to delete. Also fixed formatting. [Doug Campbell]

1.12.202 v3.1.1 (2016-10-20)

- Updated lambda policy after AWS addressed a bug in their documentation: <http://docs.aws.amazon.com/lambda/latest/dg/policy-templates.html#LambdaVPCAccessExecutionRole>. [Joel Vasallo]
- Fixed line lengths to be 120 chars for import statements using isort. [Joel Vasallo]

1.12.203 v3.1.0 (2016-10-12)

Fix

- No longer need to split the logging string. [Nate Tangsurat]

Other

- Fixed logging and docstrings to match convention. [Joel Vasallo]

- Added log messages and fixed default tls_naming.json.j2 path. [Joel Vasallo]

- Fixed import sorts and standardized log formatting. [Joel Vasallo]

- Created ForemastTemplateNotFound exception and updated references to use this. [Joel Vasallo]

- Added custom tls cert naming generation. [Joel Vasallo]

– This will help out rotating IAM TLS/SSL certs without impact and forcing developers to update the name in their repo files

- Raising exception if template file can't be found. [Joel Vasallo]

1.12.204 v3.0.2 (2016-10-10)

Fix

- Update to new name FileLookup. [Nate Tangsurat]
See also: #72
- Missing character. [Nate Tangsurat]
See also: #72
- CRITICAL when missing app.json files. [Nate Tangsurat]
See also: #72
- Add runway directory path to INFO. [Nate Tangsurat]
See also: #72

Other

- Refactor: Squash assignments. [Nate Tangsurat]
See also: #72
- Refactor: Remove extra INFO. [Nate Tangsurat]
See also: #72
- Refactor: Remove unused globals and imports. [Nate Tangsurat]
See also: #72
- Refactor: process_runway_configs() use GitLookup. [Nate Tangsurat]
See also: #72
- Refactor: process_git_configs() use GitLookup. [Nate Tangsurat]
See also: #72
- Style: isort, YAPF. [Nate Tangsurat]
See also: #72

1.12.205 v3.0.1 (2016-10-10)

Features

- Elasticache IAM Template. [Aaron Rea]

1.12.206 v3.0 (2016-10-10)

Fix

- Command typo. [Nate Tangsurat]

Other

- Added runway dir path to README to avoid future confussion. [Doug Campbell]
- Changed continue to pass. [Doug Campbell]
- Updated docstrings as per #81 review. [Doug Campbell]
- Updated wording to user master, branch, merge instead of a and b. [Doug Campbell]
- Updated templates to use a/b reenumerate logic. [Doug Campbell]
- Updated reenumerate function to remove hardcoded stage names. [Doug Campbell]

1.12.207 v2.20.1 (2016-10-05)

Features

- Kinesis Firehose IAM template. [Nate Tangsurat]

See also: #70

1.12.208 v2.20 (2016-10-04)

Fix

- Rename GitLookup to FileLookup. [Nate Tangsurat]
See also: #72
- Support Python 3.4 JSON decoding for now. [Nate Tangsurat]
See also: #72
- Raise for missing remote file. [Nate Tangsurat]
See also: #72
- WARNING when local file missing. [Nate Tangsurat]
See also: #72
- No Git initialization when using local runway. [Nate Tangsurat]
See also: #72
- Use Git short name in INFO. [Nate Tangsurat]
See also: #72
- Move bytes decode to get() [Nate Tangsurat]
See also: #72

Features

- Support runway directory override. [Nate Tangsurat]
See also: #74

- Centralized Git file retrieval. [Nate Tangsurat]

See also: #72

Other

- Merge pull request #74 from gogoair/feat/central_git_lookup. [Sijis Aviles]

feat: Provide centralized Git file retrieval

- Refactor: Move log statements to methods. [Nate Tangsurat]

See also: #72

- Tests: Local file lookup should not access Git. [Nate Tangsurat]

See also: #72

- Tests: Ensure FileNotFoundError raises. [Nate Tangsurat]

See also: #72

- Refactor: Extra remote Git file lookup. [Nate Tangsurat]

See also: #72

- Refactor: Move safe assignment outside of try. [Nate Tangsurat]

See also: #72

- Refactor: Extract local file lookup. [Nate Tangsurat]

See also: #72

- Docs: Update GitLookup docstring. [Nate Tangsurat]

See also: #72

- Docs: Add exception raise note. [Nate Tangsurat]

See also: #72

- Tests: Runway directory support. [Nate Tangsurat]

See also: #72

- Tests: Assert return types to match docstrings. [Nate Tangsurat]

See also: #72

- Tests: Separate variables for bytes and string. [Nate Tangsurat]

See also: #72

- Tests: Check invalid JSON from Git. [Nate Tangsurat]

See also: #72

- Tests: Check GitLookup methods. [Nate Tangsurat]

See also: #72

- Tests: GitLookup basic case. [Nate Tangsurat]

See also: #72

- Fixed a typo in s3 docs. [Joel Vasallo]

- Updated documentation around s3 buckets. [Joel Vasallo]

- Mentioned you can also provide a list of bucket names

1.12.209 v2.19.4 (2016-10-04)

Fix

- Certs need default of empty string. [Nate Tangsurat]
- Expand certificate paths. [Nate Tangsurat]
- Forgot to add lint check to Travis. [Nate Tangsurat]

Features

- ElasticSearch IAM template. [Aaron Rea]

Other

- Merge pull request #79 from gogoair/fix/expand_cert_paths. [Joel Vasallo]
fix: Expand certificate paths
- Style: isort, YAPF. [Nate Tangsurat]
See also: #72
- Split tox into code / lint checks. [Sijis Aviles]

1.12.210 v2.19.3 (2016-09-30)

Features

- Define Pipeline types allowed. [Nate Tangsurat]
See also: #72

Other

- Tests: Check default Pipeline types. [Nate Tangsurat]
See also: #73
- Docs: Pipeline types in foremast.cfg. [Nate Tangsurat]
See also: #72
- Refactor: Save Pipeline type. [Nate Tangsurat]
See also: #72

1.12.211 v2.19.2 (2016-09-30)

Fix

- Default max ASGs for redblack deployments. [Nate Tangsurat]

See also: #65

Other

- Set rebuild_pipelines failure as warnings. [Sijis Aviles]
- Fixed app typo in rebuilding all. [Doug Campbell]
- Merge pull request #69 from gogoair/feature/lambda-logging-template. [Sibin Arsenijevic]
Added lambda IAM permission to log by default
- Added lambda IAM permission to log by default. [Sibin Arsenijevic]
- Newlogo with correct text. [Doug Campbell]
- Changed to logo without text until Typo is fixed. [Doug Campbell]
- Added Logo to README. [Doug Campbell]

1.12.212 v2.19.1 (2016-09-26)

- Merge pull request #64 from gogoair/fix/lambda-trigger-permissions. [Sibin Arsenijevic]
- Removed unused imports. [Sibin Arsenijevic]
- FIX: lambda trigger permissions now targeting function alias instead of app_name. [Sibin Arsenijevic]
- Updated docs for enable_public_ips. [Joel Vasallo]
- Added support for modifying associatePublicIpAddress. [Joel Vasallo]

1.12.213 v2.19 (2016-09-23)

- Removed backslash, reformed. [Sibin Arsenijevic]
- Fixed styling on docstrings, added License to lambda files that were missing it. Other minor cleanup of unused variables. [Sibin Arsenijevic]
- Styling, logic and error handling fixes as per requests in PR#60. [Sibin Arsenijevic]
- Get_lambda_alias tests added. [Sibin Arsenijevic]
- Redefined event triggers to use lambda alias as base for triggering. [Sibin Arsenijevic]
- New lambda alias arn getter get_lambda_alias_arn() and exception to follow. [Sibin Arsenijevic]
- Created support for lambda alias-ing. Code updates lambda pointer even if the alias exists to enforce \$LATEST. [Sibin Arsenijevic]
- Merge pull request #61 from gogoair/enhancement/blog-readme. [Joel Vasallo]
added blog post to read me. Will update with future posts
- Added blog post to read me. Will update with future posts. [Doug Campbell]

- Merge pull request #57 from gogoair/bug/empty_lambda_triggers. [Sijis Aviles]
added lambda_triggers as an empty list
- Added lambda_triggers as an empty list. [Doug Campbell]

1.12.214 v2.18.1 (2016-09-15)

Fix

- Handle template defined Provider Health Check. [Nate Tangsurat]
Forgot to force enable the default Provider Health Check when the templates have this value defined.
See also: PSOBAT-2365
- Add retry to Lambda creation. [Nate Tangsurat]
Need to compensate for lag in IAM Role Policy changes.
See also: #50
- Add Sid to all service templates. [Nate Tangsurat]
See also: #51
- Add lambda to services for lambda types. [Nate Tangsurat]
When `{"type": "lambda"}` is defined, the service should automatically be added to enabled services. This will add the correct IAM Policy.
See also: #51
- Whoops, unbreak Sphinx. [Nate Tangsurat]
- PyPI doesn't like :caption:. [Nate Tangsurat]

Other

- Merge pull request #56 from gogoair/fix/provider_health_checks. [Joel Vasallo]
fix: Handle template defined Provider Health Check
- Merge pull request #53 from gogoair/fix/retry_lambda_create. [Sijis Aviles]
fix: Add retry to Lambda creation
- Merge pull request #54 from gogoair/fix/lambda_type_plus_service. [Sijis Aviles]
Fix/lambda type plus service
- Test: IAM tests should use template as base. [Nate Tangsurat]
See also: #51
- Docs: Reflow text. [Nate Tangsurat]
- Docs: Use backticks around links. [Nate Tangsurat]

1.12.215 v2.18 (2016-09-13)

Fix

- Tests need to use deepcopy. [Nate Tangsurat]

See also: PSOBAT-2465

- Rename test. [Nate Tangsurat]

See also: PSOBAT-2465

- Can't set namedtuple attributes, immutable. [Nate Tangsurat]

See also: PSOBAT-2465

- Use returned namedtuple attributes. [Nate Tangsurat]

See also: PSOBAT-2465

- Use namedtuple for returning provider HCs. [Nate Tangsurat]

See also: PSOBAT-2465

- Use provider HC when Eureka enabled. [Nate Tangsurat]

Resolves a FIXME that has needed to be done for now. Now that scaffolding is in place, use that to set the option.

See also: PSOBAT-2465

- Fallback *git describe* string for tox. [Nate Tangsurat]

When *tox* runs, *git describe* fails to generate the version number. This provides a safe default in the case when the command fails to generate a version number.

- DEBUG message wording. [Nate Tangsurat]

- Remove no Accounts check. [Nate Tangsurat]

- Wrong string format. [Nate Tangsurat]

- Raise error when no Accounts match. [Nate Tangsurat]

- DEBUG Accounts configured in Spinnaker. [Nate Tangsurat]

- Handle Spinnaker missing Accounts. [Nate Tangsurat]

Error out early when no Accounts found in Spinnaker.

Features

- Use Provider HC for Eureka enabled apps. [Nate Tangsurat]

Set the Health Check to true for the default Provider when Eureka is enabled.

See also: PSOBAT-2465

- Default to alpha releases. [Nate Tangsurat]

Other

- Merge pull request #47 from gogoair/fix/asg_healthcheck. [Sijis Aviles]
fix: Use provider HC when Eureka enabled
- Test: One provider defined for HC. [Nate Tangsurat]
See also: PSOBAT-2465
- Docs: Eureka enabled uses default HC provider. [Nate Tangsurat]
See also: PSOBAT-2465
- Test: Provider and Eureka enabled. [Nate Tangsurat]
See also: PSOBAT-2465
- Test: Eureka enabled sets Amazon HC. [Nate Tangsurat]
See also: PSOBAT-2465
- Test: Update Provider Health Check sanity. [Nate Tangsurat]
See also: PSOBAT-2465
- Test: Provider Health Check sanity. [Nate Tangsurat]
See also: PSOBAT-2465
- Refactor: Reference variable for providers. [Nate Tangsurat]
See also: PSOBAT-2465
- Refactor: Move return variables to top. [Nate Tangsurat]
See also: PSOBAT-2465
- Refactor: Extract provider health check section. [Nate Tangsurat]
See also: PSOBAT-2465
- Add travis -> gitter integration. [Sijis Aviles]
- Fix test cases as default path of api gateway resource changed. [Sijis Aviles]
- Added travis badget. [Doug Campbell]
- Merge pull request #48 from gogoair/fix/tox_setup. [Sijis Aviles]
fix: Fallback *git describe* string for tox
- Add travis-ci integration. [Sijis Aviles]
- Style: isort, YAPF. [Nate Tangsurat]
See also: PSOBAT-2465
- Merge pull request #46 from gitter-badger/gitter-badge. [Doug Campbell]
Add a Gitter chat badge to README.rst
- Add Gitter badge. [The Gitter Badger]
- Updated URLs in the readme with correct readthedocs stuff. [Doug Campbell]
- Removed sphinx.ext.githubpages as it was unnecessary and breaking readthedocs. [Doug Campbell]
- Added info around x509 certs and CA Bundle. [Joel Vasallo]
 - made a mention that only PEM is supported with both key and cert

- cleaned up wording around custom CA to make it clear for when it can be used
- added recommendation on how to use ca_bundle leveraging an existing ca bundle.
- Added x509 docs - fixes #40. [Doug Campbell]
- Docs: Use caption option for code block. [Nate Tangsurat]
- Style: Strip. [Nate Tangsurat]
- Added feedback from #44, moved testing to contributing. [Doug Campbell]
- Moved links around in Readme, added placeholder links. [Doug Campbell]
- Updated readme - new header structure. [Doug Campbell]
- Updated README by adding a list of features. [Doug Campbell]
- Updated readme, still may need some work. [Doug Campbell]
- Give the module page a better title. [Sijis Aviles]
- Reorganize sidebar for better content flow. [Sijis Aviles]
- Split Getting Started into multiple parts. [Sijis Aviles]

We need to have a quick start guide and an advance usage guide. Having them both on the same page was confusing and there was also some content overlap.

- Give releasing page better title. [Sijis Aviles]
- Add missing step when using git. [Sijis Aviles]
- Reformat code blocks. [Sijis Aviles]
- Fix spelling mistakes. [Sijis Aviles]
- Add a quick start section. [Sijis Aviles]
- Merge pull request #41 from gogoit/fix/missing_accounts. [Sijis Aviles]
fix: Handle Spinnaker missing Accounts
- Discovered duplicate hardcoded environments. Removed this. [Doug Campbell]
- Docs: Fix duplicate eureka_enabled option. [Nate Tangsurat]
Found eureka_enabled in application.json, removing false pipeline.json reference.
- Docs: Include eureka_enabled option. [Nate Tangsurat]
- Merge pull request #39 from gogoit/feat/alpha_releases. [Sijis Aviles]
feat: Default to alpha releases
- Merge pull request #37 from gogoit/rename_endpoints. [Sijis Aviles]
added foremast-* endpoints for runner
- Added in feedback from #37. [Doug Campbell]
- Changed name of onetime-pipeline to pipeline-onetime. [Doug Campbell]
- Updated docs with new endpoints. [Doug Campbell]
- Added foremast-* endpoints for runner. [Doug Campbell]
- Merge pull request #35 from gogoit/feat/auto_version. [Sijis Aviles]
feat: Get version from Tag automatically

- Docs: Update releasing workflow. [Nate Tangsurat]

1.12.216 v2.17 (2016-09-02)

Fix

- Only track annotated Tags. [Nate Tangsurat]

See also: #35

Features

- Get version from Tag automatically. [Nate Tangsurat]

Other

- Chore: Add TODO for cleaning out GitLab AMI lookup. [Nate Tangsurat]

1.12.217 v2.16.3 (2016-09-01)

Fix

- Move shebang to top. [Nate Tangsurat]

See also: PSOBAT-2503

Other

- Chore: v2.16.3. [Nate Tangsurat]
- Added slash remove. [Doug Campbell]

1.12.218 v2.16.2 (2016-08-31)

Fix

- Retry when updating Lambda Function. [Nate Tangsurat]

IAM permissions can lag a little.

See also: PSOBAT-2503

- Add ec2>CreateNetworkInterface for VPC Lambda. [Nate Tangsurat]

See also: PSOBAT-2503

Other

- Chore: v2.16.2. [Nate Tangsurat]
See also: PSOBAT-2503
- Refactor: Remove unused imports. [Nate Tangsurat]
See also: PSOBAT-2503

1.12.219 v2.16.1 (2016-08-31)

Fix

- Lambda Functions need a Spinnaker Application. [Nate Tangsurat]
See also: PSOBAT-2503
- Handle Lambda Role missing permission. [Nate Tangsurat]
See also: PSOBAT-2503

Features

- Log out Lambda Function Role ARN. [Nate Tangsurat]
- Extra logging around Lambda update. [Nate Tangsurat]

Other

- Chore: v2.16.1. [Nate Tangsurat]
See also: PSOBAT-2503
- Style: isort, YAPF. [Nate Tangsurat]

1.12.220 v2.16 (2016-08-30)

Fix

- Support default bucket in example. [Nate Tangsurat]
See also: PSOBAT-2448
- Update S3 IAM Policy. [Nate Tangsurat]
See also: PSOBAT-2448

Other

- Chore: v2.16. [Nate Tangsurat]
See also: PSOBAT-2484
- Test: Check S3 IAM Policy. [Nate Tangsurat]
See also: PSOBAT-2448

1.12.221 v2.15 (2016-08-30)

- Upped version number to 2.15. [Sibin Arsenijevic]
- Fixed logic where ELB was created even if Eureka was enabled. [Sibin Arsenijevic]

1.12.222 v2.14 (2016-08-29)

- Updated tag to 2.14. [Doug Campbell]
- Merge branch ‘master’ of [github.com:gogoit/foremast](https://github.com/gogoit/foremast). [Doug Campbell]
- Hotfix: moved lambda creation under conditional. [Doug Campbell]

1.12.223 v2.13 (2016-08-29)

Fix

- Remove default Service. [Nate Tangsurat]
See also: #32
- Remove extra keyword argument. [Nate Tangsurat]
Deployment type is already stored in *pipeline_settings*.
See also: #32

Features

- Add CloudWatch Logs to Lambda deployments. [Nate Tangsurat]
Lambda deployments should have CloudWatch Logs access automatically added.
See also: PSOBAT-2482
- Add CloudWatch Logs template. [Nate Tangsurat]
See also: PSOBAT-2482
- Pass deployment type to Policy constructor. [Nate Tangsurat]
See also: PSOBAT-2482

Other

- Chore: v2.13. [Nate Tangsurat]
- Test: Check cloudwatchlogs specifically. [Nate Tangsurat]
See also: #32
- Docs: Add cloudwatchlogs to pipeline.json services. [Nate Tangsurat]
See also: PSOBAT-2482
- Test: Check CloudWatch Logs acces for Lambda. [Nate Tangsurat]
Lambda Policy should automatically have CloudWatch Logs access.
See also: PSOBAT-2482

- Test: add_lambda_permissions moved to utils. [Nate Tangsurat]
Update testing to target new util function moved from a class method.
See also: PSOBAT-2337
- Test: Check add_lambda_permissions. [Nate Tangsurat]
See also: PSOBAT-2337
- Test: Add basic APIGateway init. [Nate Tangsurat]
See also: PSOBAT-2337
- Test: Add required rules dictionary. [Nate Tangsurat]
See also: PSOBAT-2337
- Test: API Gateway create resource. [Nate Tangsurat]
See also: PSOBAT-2337
- Added default parameter to dictionary get method. [Steven Basgall]
- Added json input passing to cloudwatch event config. [Steven Basgall]
- Added generic lambda templates. [Doug Campbell]
- Docs: Add note on documentation file extensions. [Nate Tangsurat]
See also: PSOBAT-2476
- Docs: Split documentation notes into bullets. [Nate Tangsurat]
See also: PSOBAT-2476
- Docs: Move documentation section up. [Nate Tangsurat]
See also: PSOBAT-2476
- Docs: Add contents to contributing. [Nate Tangsurat]
See also: PSOBAT-2476
- Docs: Refactor pipeline.json notifications block. [Nate Tangsurat]
- Docs: Refactor pipeline.json image block. [Nate Tangsurat]
- Docs: Change include files to .rest suffix. [Nate Tangsurat]
- Docs: Refactor pipeline.json lambda block. [Nate Tangsurat]
- Docs: Fix JSON block. [Nate Tangsurat]
- Docs: Services for pipeline.json. [Nate Tangsurat]
- Docs: Move Lambda docs into directory. [Nate Tangsurat]
Updated :doc: links to :ref:..
- Docs: Move configuration files into directory. [Nate Tangsurat]
Updated :doc: links to :ref: links.
- Docs: Move pipeline.json into directory. [Nate Tangsurat]
- Style: isort, YAPF. [Nate Tangsurat]
See also: PSOBAT-2482
- Style: isort, YAPF. [Nate Tangsurat]

- Docs: Add Branch and Commit guidelines. [Nate Tangsurat]
- Added lambda support to prod pipeline. [Fadi Almasri]
- Docs: Start contribution guidelines. [Nate Tangsurat]
- Removed References to Gate Class (in favor of requests) [Joel Vasallo]
 - We still have some work to do on this class, so I opted to standardize the way we call Gate to just use requests.
- Formatting and Missing SSL Call. [Joel Vasallo]
 - Cleaned up lines to meet 120 char limit
 - Fixed missing request.delete() call not passing SSL CA bundle and PEM
- Updated all references of requests to use x509 certificates if specified. [Joel Vasallo]
- Docs: Use wheels for releasing. [Nate Tangsurat]

1.12.224 v2.12 (2016-08-26)

Features

- Add debug flag support for runners. [Nate Tangsurat]

Other

- Chore: v2.12. [Nate Tangsurat]

1.12.225 v2.11 (2016-08-26)

Fix

- Misplaced variable. [Nate Tangsurat]

Features

- Include short Git name in INFO. [Nate Tangsurat]

Other

- Chore: v2.11. [Nate Tangsurat]
- Updated docs with feedback and editing notes from #26, added ref links. [Doug Campbell]
- Documented lambda pipeline and added image. [Doug Campbell]
- Updated lambda events docs for all values. [Doug Campbell]
- Added lambda events skeleton and aws lambda subsection. [Doug Campbell]
- Updated configuration docs. [Doug Campbell]
- Regenerated foremast code docs. Started lambda documentation. [Doug Campbell]
- Docs: Rework pipeline.json description. [Nate Tangsurat]

- Docs: Fix reference. [Nate Tangsurat]
 - Ran isort, added docstrings. [Doug Campbell]
 - Added lambda permisisons to cloudwatch log events. [Doug Campbell]
 - Updated docstrings, ran isort, feedback from #24. [Doug Campbell]
 - Added lambda permissions for SNS event. [Doug Campbell]
 - Ran isort, incorporated feedback from #23. [Doug Campbell]
 - Added lambda permissions for cloudwatch events. [Doug Campbell]
 - Added feedback from #22, changed permissions to kwargs. [Doug Campbell]
 - Added env and region to api gateway permissions. [Doug Campbell]
 - Removed print. [Doug Campbell]
 - Added lambda function util to API gateway, added to s3 creation. Intergrated all together. [Doug Campbell]
 - Updated utils to use awslambda for arn and permissions. [Doug Campbell]
 - Added additional variables for templates. [Steven Basgall]
 - Docs: Fix example configuration path. [Nate Tangsurat]
 - Docs: Remove unneeded autodoc mocks. [Nate Tangsurat]
 - Docs: Remove page for missing module. [Nate Tangsurat]
 - Docs: Fix section markers. [Nate Tangsurat]
 - Docs: Sync sphinx-apidoc. [Nate Tangsurat]
 - Docs: Move comment into docstring. [Nate Tangsurat]
 - Docs: Update docstrings. [Nate Tangsurat]
 - Docs: End docstring sentence. [Nate Tangsurat]
 - Docs: Update main Lambda docstring. [Nate Tangsurat]
 - Refactor: Remove adhoc launcher. [Nate Tangsurat]
 - Docs: Fix docstrings. [Nate Tangsurat]
 - Added feedback from #18. Refactored adding permissions and finding api and resource IDs. [Doug Campbell]
 - Removed setup api file. [Doug Campbell]
 - Added remove permissions, removed source SRN (which fixed permissions bug), ran yapf. [Doug Campbell]
 - Added doc strings, tied fuctions together for api creation. [Doug Campbell]
 - Moved functions into api_gateway. [Doug Campbell]
 - Remove unused variable. [Sijis Aviles]
 - Added doc strings, tied fuctions together for api creation. [Doug Campbell]
 - Moved functions into api_gateway. [Doug Campbell]
 - Added functions for handling API creation, need to clean up and tie everything together. [Doug Campbell]
 - Merge pull request #20 from gogoit/feature/lambda_iam. [Joel Vasallo]
- feat: Add Lamda IAM Role Trust Policy

- Refactor: Rename IAM Trust Policies. [Nate Tangsurat]
See also: PSOBAT-2092
- Test: Check Lambda IAM Role Policy. [Nate Tangsurat]
See also: PSOBAT-2092
- Chore: Add Lambda IAM Role Policy. [Nate Tangsurat]
See also: PSOBAT-2092
- Test: Check EC2 IAM Policy. [Nate Tangsurat]
See also: PSOBAT-2092
- Chore: Move IAM Role template. [Nate Tangsurat]
See also: PSOBAT-2092
- Refactor: Extract Policy template. [Nate Tangsurat]
See also: PSOBAT-2092
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-2092
- Changed default lambda runtime to java8. [Steven Basgall]
- Test: Fix subnet checking. [Nate Tangsurat]
See also: PSOBAT-2092
- Safe class selection for create_pipeline function. [Steven Basgall]
- Indentation fixes for 120 cols. [Steven Basgall]
- Changed dict to string. [Steven Basgall]

1.12.226 v2.10 (2016-08-09)

- Updated tag to 2.10. [Doug Campbell]
- Linting fixes. [Sijis Aviles]
- Updated docs. [Doug Campbell]
- Fixed elb sg issue and removed sg from configs. [Doug Campbell]
- Fixed elb sg issue and removed sg from configs. [Doug Campbell]
- Chore: Include tox for development. [Nate Tangsurat]
- Chore: Add Python 3.6 testing. [Nate Tangsurat]
- Linting fixes. [Sijis Aviles]

1.12.227 v2.9.0 (2017-05-31)

Fix

- Mock DOMAIN constant. [Nate Tangsurat]

Other

- Chore: v2.9. [Nate Tangsurat]
- Merge branch ‘master’ of github.com:gogoair/foremast. [Doug Campbell]
- Test: Fix external DNS Zone IDs mock. [Nate Tangsurat]
- Test: Fix mocking. [Nate Tangsurat]
- Test: Mock constant and request call. [Nate Tangsurat]
- Test: Fix constant mock. [Nate Tangsurat]
- Test: Fix constant mocking. [Nate Tangsurat]
- Linting fixes. [Sijis Aviles]
- Added link to foremast-templates. [Doug Campbell]
- Removed jira. [Doug Campbell]
- Updated pipeline tests. [Sijis Aviles]
- Add dns test cases. [Sijis Aviles]
- Simplified iam tests cases to just validate json. [Sijis Aviles]

It would make sense to validate the exact policy, however, policies do change significantly and the only real changes are the content.

We also cannot determine if the policy generated is in fact legal or represents what policies we’d like to enforce. Since that’s the case, i think we just ensure that generated policy is simply valid json.

- Fix elb test cases. [Sijis Aviles]
- Refactor subnets util and remove unused parameters. [Sijis Aviles]
- Renamed files in utils for easier testing. [Sijis Aviles]
- Add additional subnets tests cases. [Sijis Aviles]
- Add utils/subnets test cases. [Sijis Aviles]
- Add test cases for utils/vpc. [Sijis Aviles]
- Add utils/security_group test cases. [Sijis Aviles]
- Add utils/dns test cases. [Sijis Aviles]
- Add application util related test cases. [Sijis Aviles]
- Added more pipeline util test cases. [Sijis Aviles]
- Add simple test case for utils/slack. [Sijis Aviles]
- Add test cases for find_elb. [Sijis Aviles]
- Consolidated filename test into utils. [Sijis Aviles]
- Add test cases for checking managed pipeline. [Sijis Aviles]
- Add deep_chain_map test cases. [Sijis Aviles]
- Add utils.banners test case. [Sijis Aviles]
- Rename utils.banner -> utils.banners. [Sijis Aviles]
- Stupid print statement. [Doug Campbell]

- Moved default security groups to a config setting. [Doug Campbell]
- Removed print statement, was overlooked. [Doug Campbell]
- Updated securitygroups to not include apps-all or offices_all by default. [Doug Campbell]
- Removed gogo-archaius references. [Doug Campbell]
- Made templates more generic for open source users. [Doug Campbell]

1.12.228 v2.9 (2016-08-08)

Fix

- Mock DOMAIN constant. [Nate Tangsurat]

Other

- Chore: v2.9. [Nate Tangsurat]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Doug Campbell]
- Test: Fix external DNS Zone IDs mock. [Nate Tangsurat]
- Test: Fix mocking. [Nate Tangsurat]
- Test: Mock constant and request call. [Nate Tangsurat]
- Test: Fix constant mock. [Nate Tangsurat]
- Test: Fix constant mocking. [Nate Tangsurat]
- Linting fixes. [Sijis Aviles]
- Added link to foremast-templates. [Doug Campbell]
- Removed jira. [Doug Campbell]
- Updated pipeline tests. [Sijis Aviles]
- Add dns test cases. [Sijis Aviles]
- Simplified iam tests cases to just validate json. [Sijis Aviles]

It would make sense to validate the exact policy, however, policies do change significantly and the only real changes are the content.

We also cannot determine if the policy generated is in fact legal or represents what policies we’d like to enforce. Since that’s the case, i think we just ensure that generated policy is simply valid json.

- Fix elb test cases. [Sijis Aviles]
- Refactor subnets util and remove unused parameters. [Sijis Aviles]
- Renamed files in utils for easier testing. [Sijis Aviles]
- Add additional subnets tests cases. [Sijis Aviles]
- Add utils/subnets test cases. [Sijis Aviles]
- Add test cases for utils/vpc. [Sijis Aviles]
- Add utils/security_group test cases. [Sijis Aviles]
- Add utils/dns test cases. [Sijis Aviles]

- Add application util related test cases. [Sijis Aviles]
- Added more pipeline util test cases. [Sijis Aviles]
- Add simple test case for utils/slack. [Sijis Aviles]
- Add test cases for find_elb. [Sijis Aviles]
- Consolidated filename test into utils. [Sijis Aviles]
- Add test cases for checking managed pipeline. [Sijis Aviles]
- Add deep_chain_map test cases. [Sijis Aviles]
- Add utils.banners test case. [Sijis Aviles]
- Rename utils.banner -> utils.banners. [Sijis Aviles]
- Stupid print statement. [Doug Campbell]
- Moved default security groups to a config setting. [Doug Campbell]
- Removed print statement, was overlooked. [Doug Campbell]
- Updated securitygroups to not include apps-all or offices_all by default. [Doug Campbell]
- Removed gogo-archaius references. [Doug Campbell]
- Made templates more generic for open source users. [Doug Campbell]

1.12.229 v2.8.0 (2017-05-31)

Fix

- Generalize module to “lookups” [Nate Tangsurat]
- Move section check into value checker. [Nate Tangsurat]
- WARNING message for missing configurations. [Nate Tangsurat]

Other

- Bump to v2.8. [Sijis Aviles]
- Fixed lingering merge conflict text. [Doug Campbell]
- Added the correct pictures. [Doug Campbell]
- Fixed merge conflict. [Doug Campbell]
- Updated pipeline example and workflow doc, included pictures. [Doug Campbell]
- Added single pipeline image. [Doug Campbell]
- Added vpc subnet assumptions to docs. [Doug Campbell]
- Updated infrastructure docs. Still needs VPC info. [Doug Campbell]
- Started infra docs. [Doug Campbell]
- Fixed example references, fixed typos and basic editing, improved wording. [Doug Campbell]
- Added more content. [Doug Campbell]
- Updated doc images. [Doug Campbell]

- Updated pipeline examples for new external pipeline pathing. [Doug Campbell]
- Updated pipeline example and workflow doc, included pictures. [Doug Campbell]
- Added single pipeline image. [Doug Campbell]
- Added vpc subnet assumptions to docs. [Doug Campbell]
- Updated infrastructure docs. Still needs VPC info. [Doug Campbell]
- Started infra docs. [Doug Campbell]
- Fix bug with autoscaling policies and reorganize the code a bit. [Doug Campbell]
- Fixed merge conflict. [Doug Campbell]
- Fixed template paths in test. [Doug Campbell]
- Fixed error with expand-user. [Doug Campbell]
- Renamed template files to j2 extension. [Doug Campbell]
- Changed logic for jinja loader to accept a list. Added os.path.expanduser. [Doug Campbell]
- Merge request feedback, changed to os.path.join and removed redundant logs. [Doug Campbell]
- Simplified templatedir logic. [Doug Campbell]
- Renamed template files to remove redundant naming. [Doug Campbell]
- Renamed slack directory for templates. [Doug Campbell]
- Updated pathing on all templates. [Doug Campbell]
- Restored templates back to original, just with new pathing. [Doug Campbell]
- Added logic for template_dir config, moved pipeline-templates to just pipeline. [Doug Campbell]
- Updaated appending logic for templates. [Doug Campbell]
- Fixed template paths in test. [Doug Campbell]
- Fixed error with expand-user. [Doug Campbell]
- Renamed template files to j2 extension. [Doug Campbell]
- Changed logic for jinja loader to accept a list. Added os.path.expanduser. [Doug Campbell]
- Merge request feedback, changed to os.path.join and removed redundant logs. [Doug Campbell]
- Simplified templatedir logic. [Doug Campbell]
- Renamed template files to remove redundant naming. [Doug Campbell]
- Renamed slack directory for templates. [Doug Campbell]
- Updated pathing on all templates. [Doug Campbell]
- Restored templates back to original, just with new pathing. [Doug Campbell]
- Added logic for template_dir config, moved pipeline-templates to just pipeline. [Doug Campbell]
- Chore: Add FIXME note for settings ASG HC to EC2. [Nate Tangsurat]

Need to make sure that `provider_healthcheck` is set when `eureka_enabled`. Without it, the ASG Health Check defaults back to ELB in Spinnaker.

- Remove IDE project files. [Sijis Aviles]
- Licensing as Apache 2.0. [Sijis Aviles]

- Fixed for loop else and put app_name in a variable for gogoutils. [Doug Campbell]
- Incorporated feedback from merge request 13. [Doug Campbell]
- Updated docstrings and conditional wrapping. [Doug Campbell]
- Added functions to look up existing pipeline and compare with desired pipelines. This puts pipeline_id in the templates. [Doug Campbell]
- Updated docstrings and conditional wrapping. [Doug Campbell]
- Updated gogoutils call to remove deprecation warning. [Doug Campbell]
- Added functions to look up existing pipeline and compare with desired pipelines. This puts pipeline_id in the templates. [Doug Campbell]
- Removed runway contents to separate repo. [Sijis Aviles]
- Removed old file. [Doug Campbell]
- Merge conflict fix. [Doug Campbell]
- Added new task endpoint for elb and sg's. [Doug Campbell]
- Removed appname for checktask and updated scalingpolicy. [Doug Campbell]
- Added post_task utils and updated create_app for new endpoint. [Doug Campbell]
- Removed deprecated API calls for tasks, using just /tasks now. [Doug Campbell]
- Moved functions to tasks.py, updated response variable name. [Doug Campbell]
- Added new task endpoint for elb and sg's. [Doug Campbell]
- Removed appname for checktask and updated scalingpolicy. [Doug Campbell]
- Added post_task utils and updated create_app for new endpoint. [Doug Campbell]
- Removed deprecated API calls for tasks, using just /tasks now. [Doug Campbell]
- Revert “Remove condition for eureka flag” [Sijis Aviles]

This reverts commit 0e70e27f30e826e46ef16c72c3810f2799b3f81d.

- Change default builder to ebs. [Sijis Aviles]
- Combine prospector with existing pytest command. [Sijis Aviles]
- Enable pyflakes and disable checking of __init__.py. [Sijis Aviles]
- Use better options for lint. [Sijis Aviles]
- Add lint tox command. [Sijis Aviles]
- Run tests in tests directory. [Sijis Aviles]
- Add consts test cases. [Sijis Aviles]
- Add missing docstring. [Sijis Aviles]
- Renamed variable. [Sijis Aviles]
- Include latest gogoutils feature that includes formatting. [Sijis Aviles]
- Moved cleanup into the try/except. [Doug Campbell]
- Fixed typo. [Doug Campbell]
- Added try catch around rebuilding. [Doug Campbell]
- Docs: Sync modules. [Nate Tangsurat]

- Test: Start testing around ami_lookup() [Nate Tangsurat]

Need to rename *foremast.utils.ami_lookup* the module. Mocking can't target the module when *ami_lookup()* is masking the path.

- Update order of when config file are read. [Sijis Aviles]

Configuration files at lower levels (in user directories) should override any values set a global or higher level.

- Test: Enable pep8 and enforce line length 120. [Nate Tangsurat]

- Fixed a typo in servicenow stage. [Fadi Almasri]

- Merge branch 'feature/configs' into 'master' [Nate Tangsurat]

fix: WARNING message for missing configurations

Removes the *raise SystemExit* in favour of warning messages when sourcing *foremast.consts*. This will allow tests to run with a missing configuration file and mocking should be possible.

See merge request !7

- Merge branch 'feature/chroot_options' [Sijis Aviles]
- Use imperative sentences. [Sijis Aviles]
- Add feature to determine how baking is handled. [Sijis Aviles]
- Remove token, as its handled by new config format. [Sijis Aviles]
- Give tests better names. [Sijis Aviles]
- Merge branch 'cross-account-sg' into 'master' [Sijis Aviles]

Enabled Cross Account VPC Security Groups

Doing this should enable cross account security group access.

My cases I should handle: - Developer specifies app name SG, specifies account, specifies ports. - Developer doesn't specify app name, account, or ports, - Developer specifies app_name, ports, but not account.

Logic: If account is specified, cross account is true. Look up vpcID and post to spinnaker to enable cross account SG. If account is not specified, default to legacy behavior.

Here is the sample JSON posted to Spinnaker. **“““javascript**

```
{ "accountName": "dev", "id": "sg-578bcc2f", "name": "coreforrest", "type": "tcp",
  "startPort": 80, "endPort": 80, "existing": true
}, {
  "type": "tcp", "startPort": 7001, "endPort": 7001, "name": "admincrewconnect",
  "crossAccountEnabled": true, "accountName": "dev", "vpcId": "vpc-869f46e2"
}
],  
“““
```

See merge request !5

- Mock actual update calls when running security group tests. [Sijis Aviles]
- Added testcases for security group. [Joel Vasallo]
- Renamed key from account to env since it will make it clearer for developers to understand. [Joel Vasallo]
- Cleaned up cross account SG feature. [Joel Vasallo]

- WIP: Fixed cross account SG to actually work. [Joel Vasallo]
- Enabled Cross Account VPC Security Groups. [Joel Vasallo]
- Updated readme, made better for general public. [Doug Campbell]
- Fixed missing key error for policies in legacy ELB way. [Joel Vasallo]
- Updated getting_started guide. [Doug Campbell]
- Added getting started docs. [Doug Campbell]
- Added doc for aws credentials. [Doug Campbell]
- Merge branch ‘feature/dir’ of github.com:gogoair/foremast into feature/dir. [Doug Campbell]
- Fixed potential bugs with empty lists, none evalutions, and logging. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Merge branch ‘feature/dir’ of github.com:gogoair/foremast into feature/dir. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Merge branch ‘master’ of github.com:gogoair/foremast into feature/dir. [Doug Campbell]
- Docs: Clean up foremast.configs docstring. [Nate Tangsurat]
- Chore: Ignore generated files. [Nate Tangsurat]
- Docs: Forgot to include link to requirements. [Nate Tangsurat]
- Docs: Change configuration example to link. [Nate Tangsurat]
- Docs: Remove extra underline for links. [Nate Tangsurat]
- Docs: Use code-block. [Nate Tangsurat]
- Docs: Use code-block. [Nate Tangsurat]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]

- Added better logging and changed dict to .get() for configs. [Doug Campbell]
- Fixed token error. [Doug Campbell]
- Added runway_dir env. [Doug Campbell]
- Fixed potential bugs with empty lists, none evalutions, and logging. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Added better logging and changed dict to .get() for configs. [Doug Campbell]
- Fixed token error. [Doug Campbell]
- Added runway_dir env. [Doug Campbell]

1.12.230 v2.8 (2016-08-08)

Fix

- Generalize module to “lookups” [Nate Tangsurat]
- Move section check into value checker. [Nate Tangsurat]
- WARNING message for missing configurations. [Nate Tangsurat]

Other

- Bump to v2.8. [Sijis Aviles]
- Fixed lingering merge conflict text. [Doug Campbell]
- Added the correct pictures. [Doug Campbell]
- Fixed merge conflict. [Doug Campbell]
- Updated pipeline example and workflow doc, included pictures. [Doug Campbell]
- Added single pipeline image. [Doug Campbell]
- Added vpc subnet assumptions to docs. [Doug Campbell]
- Updated infrastructure docs. Still needs VPC info. [Doug Campbell]
- Started infra docs. [Doug Campbell]
- Fixed example references, fixed typos and basic editing, improved wording. [Doug Campbell]
- Added more content. [Doug Campbell]
- Updated doc images. [Doug Campbell]
- Updated pipeline examples for new external pipeline pathing. [Doug Campbell]

- Updated pipeline example and workflow doc, included pictures. [Doug Campbell]
- Added single pipeline image. [Doug Campbell]
- Added vpc subnet assumptions to docs. [Doug Campbell]
- Updated infrastructure docs. Still needs VPC info. [Doug Campbell]
- Started infra docs. [Doug Campbell]
- Fix bug with autoscaling policies and reorganize the code a bit. [Doug Campbell]
- Fixed merge conflict. [Doug Campbell]
- Fixed template paths in test. [Doug Campbell]
- Fixed error with expand-user. [Doug Campbell]
- Renamed template files to j2 extension. [Doug Campbell]
- Changed logic for jinja loader to accept a list. Added os.path.expanduser. [Doug Campbell]
- Merge request feedback, changed to os.path.join and removed redundant logs. [Doug Campbell]
- Simplified templatedir logic. [Doug Campbell]
- Renamed template files to remove redundant naming. [Doug Campbell]
- Renamed slack directory for templates. [Doug Campbell]
- Updated pathing on all templates. [Doug Campbell]
- Restored templates back to original, just with new pathing. [Doug Campbell]
- Added logic for template_dir config, moved pipeline-templates to just pipeline. [Doug Campbell]
- Updated appending logic for templates. [Doug Campbell]
- Fixed template paths in test. [Doug Campbell]
- Fixed error with expand-user. [Doug Campbell]
- Renamed template files to j2 extension. [Doug Campbell]
- Changed logic for jinja loader to accept a list. Added os.path.expanduser. [Doug Campbell]
- Merge request feedback, changed to os.path.join and removed redundant logs. [Doug Campbell]
- Simplified templatedir logic. [Doug Campbell]
- Renamed template files to remove redundant naming. [Doug Campbell]
- Renamed slack directory for templates. [Doug Campbell]
- Updated pathing on all templates. [Doug Campbell]
- Restored templates back to original, just with new pathing. [Doug Campbell]
- Added logic for template_dir config, moved pipeline-templates to just pipeline. [Doug Campbell]
- Chore: Add FIXME note for settings ASG HC to EC2. [Nate Tangsurat]

Need to make sure that `provider_healthcheck` is set when `eureka_enabled`. Without it, the ASG Health Check defaults back to ELB in Spinnaker.

- Remove IDE project files. [Sijis Aviles]
- Licensing as Apache 2.0. [Sijis Aviles]
- Fixed for loop else and put app_name in a variable for gogoutils. [Doug Campbell]

- Incorporated feedback from merge request 13. [Doug Campbell]
- Updated docstrings and conditional wrapping. [Doug Campbell]
- Added functions to look up existing pipeline and compare with desired pipelines. This puts pipeline_id in the templates. [Doug Campbell]
- Updated docstrings and conditional wrapping. [Doug Campbell]
- Updated gogoutils call to remove deprecation warning. [Doug Campbell]
- Added functions to look up existing pipeline and compare with desired pipelines. This puts pipeline_id in the templates. [Doug Campbell]
- Removed runway contents to separate repo. [Sijis Aviles]
- Removed old file. [Doug Campbell]
- Merge conflict fix. [Doug Campbell]
- Added new task endpoint for elb and sg's. [Doug Campbell]
- Removed appname for checktask and updated scalingpolicy. [Doug Campbell]
- Added post_task utils and updated create_app for new endpoint. [Doug Campbell]
- Removed deprecated API calls for tasks, using just /tasks now. [Doug Campbell]
- Moved functions to tasks.py, updated response variable name. [Doug Campbell]
- Added new task endpoint for elb and sg's. [Doug Campbell]
- Removed appname for checktask and updated scalingpolicy. [Doug Campbell]
- Added post_task utils and updated create_app for new endpoint. [Doug Campbell]
- Removed deprecated API calls for tasks, using just /tasks now. [Doug Campbell]
- Revert “Remove condition for eureka flag” [Sijis Aviles]
This reverts commit 0e70e27f30e826e46ef16c72c3810f2799b3f81d.
- Change default builder to ebs. [Sijis Aviles]
- Combine prospector with existing pytest command. [Sijis Aviles]
- Enable pyflakes and disable checking of __init__.py. [Sijis Aviles]
- Use better options for lint. [Sijis Aviles]
- Add lint tox command. [Sijis Aviles]
- Run tests in tests directory. [Sijis Aviles]
- Add consts test cases. [Sijis Aviles]
- Add missing docstring. [Sijis Aviles]
- Renamed variable. [Sijis Aviles]
- Include latest gogoutils feature that includes formatting. [Sijis Aviles]
- Moved cleanup into the try/except. [Doug Campbell]
- Fixed typo. [Doug Campbell]
- Added try catch around rebuilding. [Doug Campbell]
- Docs: Sync modules. [Nate Tangsurat]

- Test: Start testing around ami_lookup() [Nate Tangsurat]

Need to rename *foremast.utils.ami_lookup* the module. Mocking can't target the module when *ami_lookup()* is masking the path.

- Update order of when config file are read. [Sijis Aviles]

Configuration files at lower levels (in user directories) should override any values set a global or higher level.

- Test: Enable pep8 and enforce line length 120. [Nate Tangsurat]

- Fixed a typo in servicenow stage. [Fadi Almasri]

- Merge branch 'feature/configs' into 'master' [Nate Tangsurat]

fix: WARNING message for missing configurations

Removes the *raise SystemExit* in favour of warning messages when sourcing *foremast.consts*. This will allow tests to run with a missing configuration file and mocking should be possible.

See merge request !7

- Merge branch 'feature/chroot_options' [Sijis Aviles]
- Use imperative sentences. [Sijis Aviles]
- Add feature to determine how baking is handled. [Sijis Aviles]
- Remove token, as its handled by new config format. [Sijis Aviles]
- Give tests better names. [Sijis Aviles]
- Merge branch 'cross-account-sg' into 'master' [Sijis Aviles]

Enabled Cross Account VPC Security Groups

Doing this should enable cross account security group access.

My cases I should handle: - Developer specifies app name SG, specifies account, specifies ports. - Developer doesn't specify app name, account, or ports, - Developer specifies app_name, ports, but not account.

Logic: If account is specified, cross account is true. Look up vpcID and post to spinnaker to enable cross account SG. If account is not specified, default to legacy behavior.

Here is the sample JSON posted to Spinnaker. **“““javascript**

```
{ "accountName": "dev", "id": "sg-578bcc2f", "name": "coreforrest", "type": "tcp",
  "startPort": 80, "endPort": 80, "existing": true
}, {
  "type": "tcp", "startPort": 7001, "endPort": 7001, "name": "admincrewconnect",
  "crossAccountEnabled": true, "accountName": "dev", "vpcId": "vpc-869f46e2"
}
],  
“““
```

See merge request !5

- Mock actual update calls when running security group tests. [Sijis Aviles]
- Added testcases for security group. [Joel Vasallo]
- Renamed key from account to env since it will make it clearer for developers to understand. [Joel Vasallo]
- Cleaned up cross account SG feature. [Joel Vasallo]

- WIP: Fixed cross account SG to actually work. [Joel Vasallo]
- Enabled Cross Account VPC Security Groups. [Joel Vasallo]
- Updated readme, made better for general public. [Doug Campbell]
- Fixed missing key error for policies in legacy ELB way. [Joel Vasallo]
- Updated getting_started guide. [Doug Campbell]
- Added getting started docs. [Doug Campbell]
- Added doc for aws credentials. [Doug Campbell]
- Merge branch ‘feature/dir’ of git.gogoair.com:spinnaker/pipes into feature/dir. [Doug Campbell]
- Fixed potential bugs with empty lists, none evalutions, and logging. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Merge branch ‘feature/dir’ of git.gogoair.com:spinnaker/pipes into feature/dir. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes into feature/dir. [Doug Campbell]
- Docs: Clean up foremast.configs docstring. [Nate Tangsurat]
- Chore: Ignore generated files. [Nate Tangsurat]
- Docs: Forgot to include link to requirements. [Nate Tangsurat]
- Docs: Change configuration example to link. [Nate Tangsurat]
- Docs: Remove extra underline for links. [Nate Tangsurat]
- Docs: Use code-block. [Nate Tangsurat]
- Docs: Use code-block. [Nate Tangsurat]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]

- Added better logging and changed dict to .get() for configs. [Doug Campbell]
- Fixed token error. [Doug Campbell]
- Added runway_dir env. [Doug Campbell]
- Fixed potential bugs with empty lists, none evalutions, and logging. [Doug Campbell]
- Docs: Punctuation. [Nate Tangsurat]
- Docs: Use code-block for JSON. [Nate Tangsurat]
- Added fall back for just ami ID and updated docs. [Doug Campbell]
- Updated docs with ami_json. [Doug Campbell]
- Added docs. [Doug Campbell]
- Changed ami-url to a base url. [Doug Campbell]
- Added ami lookup from URL, not just gitlab. [Doug Campbell]
- Added better logging and changed dict to .get() for configs. [Doug Campbell]
- Fixed token error. [Doug Campbell]
- Added runway_dir env. [Doug Campbell]

1.12.231 v2.7.0 (2017-05-31)

- Bump to v2.7. [Sijis Aviles]
- Merge branch ‘root-volume’ [Sijis Aviles]
- Expose ability to update size of root volume on image. [Sijis Aviles]
- Added support for specifying root_volume_size. [Joel Vasallo]
- Increased healthcheck grace to 3 minutes. [Doug Campbell]
- Added condition to switch HC to EC2 in case there is “eureka_enabled” in runway files. [Sibin Arsenijevic]

1.12.232 v2.7 (2016-06-24)

- Bump to v2.7. [Sijis Aviles]
- Merge branch ‘root-volume’ [Sijis Aviles]
- Expose ability to update size of root volume on image. [Sijis Aviles]
- Added support for specifying root_volume_size. [Joel Vasallo]
- Increased healthcheck grace to 3 minutes. [Doug Campbell]
- Added condition to switch HC to EC2 in case there is “eureka_enabled” in runway files. [Sibin Arsenijevic]

1.12.233 v2.6.0 (2017-05-31)

- Updated version v2.6. [Sijis Aviles]
- Remove strings so its empty. [Sijis Aviles]
- Remove redundant check to set defaults. [Sijis Aviles]

- Add message when unable to decode a json file. [Sijis Aviles]
- Fixed template formate. [Doug Campbell]
- Added owner_email field. [Jeff Lu]
- Fixed formatting issue. [Doug Campbell]
- Took out earlier changes. [Jeff Lu]
- Adds owner_email to pipeline template. [Jeff Lu]
- Added better logic for template processing. [Doug Campbell]
- Forgot comma. [Doug Campbell]
- Continuing to test templating logic. [Doug Campbell]
- Fixed issue with lenght in jinja2. [Doug Campbell]
- Added conditional in stage-deploy for interestingHealthProviderName. [Doug Campbell]
- Added documentation key to pipeline.json. [Doug Campbell]
- Docs: Copy releasing information. [Nate Tangsurat]

1.12.234 v2.6 (2016-06-22)

Fix

- Bring modules into package level. [Nate Tangsurat]
- Bring destroy functions up a few levels. [Nate Tangsurat]
- Remove trailing slash for DynamoDB Policy. [Nate Tangsurat]
- Use different exception. [Nate Tangsurat]

See also: PSOBAT-1538

- Calm the *full-destroy* logging a bit. [Nate Tangsurat]

See also: PSOBAT-1538

- Handle *destroy-sg* boto3 AccessDenied error. [Nate Tangsurat]

See also: PSOBAT-1538

- Ignore *destroy-iam* boto3 AccessDenied. [Nate Tangsurat]

See also: PSOBAT-1538

- Remove SystemExit from *destroy-sg* [Nate Tangsurat]

See also: PSOBAT-1538

- Ignore DNS AccessDenied destroy error. [Nate Tangsurat]

See also: PSOBAT-1538

- Use format string placeholder. [Nate Tangsurat]

See also: PSOBAT-1803

- Set banner logger to INFO. [Nate Tangsurat]

See also: PSOBAT-1803

Features

- Add banner painting utility. [Nate Tangsurat]
- Add *pipes-destroyer* Job. [Nate Tangsurat]
See also: PSOBAT-1538
- Add *destroy.bash* script for Jenkins Job. [Nate Tangsurat]
See also: PSOBAT-1538
- Add banner painting utility. [Nate Tangsurat]

Other

- Updated version v2.6. [Sijis Aviles]
- Remove strings so its empty. [Sijis Aviles]
- Remove redundant check to set defaults. [Sijis Aviles]
- Add message when unable to decode a json file. [Sijis Aviles]
- Fixed template formate. [Doug Campbell]
- Added owner_email field. [Jeff Lu]
- Fixed formatting issue. [Doug Campbell]
- Took out earlier changes. [Jeff Lu]
- Adds owner_email to pipeline template. [Jeff Lu]
- Added better logic for template processing. [Doug Campbell]
- Forgot comma. [Doug Campbell]
- Continuing to test templating logic. [Doug Campbell]
- Fixed issue with lenght in jinja2. [Doug Campbell]
- Added conditional in stage-deploy for interestingHealthProviderName. [Doug Campbell]
- Added documentation key to pipeline.json. [Doug Campbell]
- Docs: Copy releasing information. [Nate Tangsurat]
- V2.5. [Nate Tangsurat]
- Sending production notifications to a single channel. [Sijis Aviles]
- Changed redblack startegy to scale down to 0 instances. [Doug Campbell]
- Added appversion to inject build data into spinnaker. [Joel Vasallo]
- Forgot a comma, classic mistake. [Doug Campbell]
- Added deploy_strategy to templates. [Doug Campbell]
- Fixed merge conflix with banner. [Doug Campbell]
- Add missing trigger job variable. [Sijis Aviles]
- Update infra stage to reference proper jobname. [Sijis Aviles]
- Added a space for formating. [Doug Campbell]
- Reordered group and repo, bug. [Doug Campbell]

- Disabled slack notifications for now. [Doug Campbell]
- Updated dict to contain generated repo name. [Doug Campbell]
- Added infra template and updated pipeline templates. [Doug Campbell]
- Add missing token. [Sijis Aviles]
- Updated perms so its executable. [Sijis Aviles]
- Merge create-app and create pipeline into a single job. [Sijis Aviles]
- Add env and region parameters to prepare. [Sijis Aviles]
- Refactor: Remove unused imports. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: Update *full-destroy* docstring. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: INFO end of destruction. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: INFO destroy loop. [Nate Tangsurat]
See also: PSOBAT-1538
- Updated jenkins jobs for new python runner. [Doug Campbell]
- Docs: Add foremast package docstring. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Lazy logging. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Update docstrings. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Use lazy logging format. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: More descriptive variable name. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Strip. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803

- Style: Sort console scripts. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused variable. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Reflow TODO comment. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Remove example in favour of argparse help. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: isort. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803

1.12.235 v2.5.0 (2017-05-31)

Fix

- Bring modules into package level. [Nate Tangsurat]
- Bring destroy functions up a few levels. [Nate Tangsurat]
- Remove trailing slash for DynamoDB Policy. [Nate Tangsurat]
- Use different exception. [Nate Tangsurat]
See also: PSOBAT-1538
- Calm the *full-destroy* logging a bit. [Nate Tangsurat]
See also: PSOBAT-1538
- Handle *destroy-sg* boto3 AccessDenied error. [Nate Tangsurat]
See also: PSOBAT-1538
- Ignore *destroy-iam* boto3 AccessDenied. [Nate Tangsurat]
See also: PSOBAT-1538

- Remove SystemExit from *destroy-sg* [Nate Tangsurat]

See also: PSOBAT-1538

- Ignore DNS AccessDenied destroy error. [Nate Tangsurat]

See also: PSOBAT-1538

- Handle empty Gate response. [Nate Tangsurat]

See also: PSOBAT-1538

- Ignore Security Group delete errors. [Nate Tangsurat]

See also: PSOBAT-1538

- Ignore any SpinnakerError Exception. [Nate Tangsurat]

See also: PSOBAT-1538

- Use format string placeholder. [Nate Tangsurat]

See also: PSOBAT-1803

- Set banner logger to INFO. [Nate Tangsurat]

See also: PSOBAT-1803

- Missing comma in arguments. [Nate Tangsurat]

See also: PSOBAT-1803

- Squish string format down to one format. [Nate Tangsurat]

See also: PSOBAT-1803

- Use empty string for string arguments. [Nate Tangsurat]

See also: PSOBAT-1803

- Remove try, config defaults ensure key exists. [Nate Tangsurat]

See also: PSOBAT-1803

- Saved configs needs to be deep merged version. [Nate Tangsurat]

Returned data from *configs.process_git_configs()* is the contents of *pipeline.json* with some defaults. The full configuration after that has been overlayed on top of the default configuration template is from *configs.write_variables()*.

See also: PSOBAT-1803

Features

- Add banner painting utility. [Nate Tangsurat]

- Add *pipes-destroyer* Job. [Nate Tangsurat]

See also: PSOBAT-1538

- Add *destroy.bash* script for Jenkins Job. [Nate Tangsurat]

See also: PSOBAT-1538

- Add *full-destroy* command. [Nate Tangsurat]

See also: PSOBAT-1538

- Add destroyer script. [Nate Tangsurat]
See also: PSOBAT-1538
- Add banner painting utility. [Nate Tangsurat]

Other

- V2.5. [Nate Tangsurat]
- Sending production notifications to a single channel. [Sijis Aviles]
- Changed redblack startegy to scale down to 0 instances. [Doug Campbell]
- Added appversion to inject build data into spinnaker. [Joel Vasallo]
- Forgot a comma, classic mistake. [Doug Campbell]
- Added deploy_strategy to templates. [Doug Campbell]
- Fixed merge conflux with banner. [Doug Campbell]
- Add missing trigger job variable. [Sijis Aviles]
- Update infra stage to reference proper jobname. [Sijis Aviles]
- Added a space for formating. [Doug Campbell]
- Reordered group and repo, bug. [Doug Campbell]
- Disabled slack notifications for now. [Doug Campbell]
- Updated dict to contain generated repo name. [Doug Campbell]
- Added infra template and updated pipeline templates. [Doug Campbell]
- Add missing token. [Sijis Aviles]
- Updated perms so its executable. [Sijis Aviles]
- Merge create-app and create pipeline into a single job. [Sijis Aviles]
- Add env and region parameters to prepare. [Sijis Aviles]
- Refactor: Remove unused imports. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: Update *full-destroy* docstring. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: INFO end of destruction. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: INFO destroy loop. [Nate Tangsurat]
See also: PSOBAT-1538
- Docs: DEBUG Gate response. [Nate Tangsurat]
See also: PSOBAT-1538
- Disabled slack notification for team defined channels. This is noisy, will work on a better solution. [Doug Campbell]
- Clean up from bug fix. [Doug Campbell]

- Continued bug fix for slack post. [Doug Campbell]
- Bug fix for slack notify. [Doug Campbell]
- Bug fix for slack notify not finding channel. [Doug Campbell]
- Updated wording on slack template. [Doug Campbell]
- Updated utils to work for slack message. [Doug Campbell]
- Forgot self. [Doug Campbell]
- Added actual post message. [Doug Campbell]
- Fixed typo. [Doug Campbell]
- Updated jenkins jobs for new python runner. [Doug Campbell]
- Docs: Add foremast package docstring. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Lazy logging. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Update docstrings. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Use lazy logging format. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: More descriptive variable name. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Strip. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Sort console scripts. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Remove unused variable. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803

- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Reflow TODO comment. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Remove example in favour of argparse help. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: isort. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Use empty strings for keyword arguments. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Update docstring. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Reflow comment. [Nate Tangsurat]
See also: PSOBAT-1803
- Docs: Update docstring. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: isort, YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Move method entry points to functions. [Nate Tangsurat]
Keep the class pure operational calls and have the calling entry point functions tie the calls together.
See also: PSOBAT-1803
- Refactor: Reuse keyword arguments. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: Collapse small method lines. [Nate Tangsurat]
See also: PSOBAT-1803
- Style: YAPF. [Nate Tangsurat]
See also: PSOBAT-1803
- Refactor: Consolidate private token check. [Nate Tangsurat]
See also: PSOBAT-1803

- Docs: Init has no arguments. [Nate Tangsurat]

See also: PSOBAT-1803

1.12.236 v2.4.0 (2017-05-31)

Features

- Add *destroy-dns* command. [Nate Tangsurat]

See also: PSOBAT-1535

- Add util for getting DNS Zone IDs. [Nate Tangsurat]

See also: PSOBAT-1535

Other

- V2.4. [Nate Tangsurat]

1.12.237 v2.4 (2016-05-17)

Features

- Add *destroy-dns* command. [Nate Tangsurat]

See also: PSOBAT-1535

- Add util for getting DNS Zone IDs. [Nate Tangsurat]

See also: PSOBAT-1535

Other

- V2.4. [Nate Tangsurat]

1.12.238 v2.3.0 (2017-05-31)

Fix

- ELB arguments replaced by default configs. [Nate Tangsurat]

See also: PSOBAT-1534

Other

- V2.3. [Nate Tangsurat]

- Style: Use newlines for easier to read INFO. [Nate Tangsurat]

See also: PSOBAT-1534

1.12.239 v2.3 (2016-05-12)

Fix

- ELB arguments replaced by default configs. [Nate Tangsurat]

See also: PSOBAT-1534

Other

- V2.3. [Nate Tangsurat]
- Style: Use newlines for easier to read INFO. [Nate Tangsurat]

See also: PSOBAT-1534

1.12.240 v2.2.0 (2017-05-31)

Features

- Add *destroy-elb* command. [Nate Tangsurat]

See also: PSOBAT-1534

Other

- V2.2. [Nate Tangsurat]
 - Chore: Remove unused comment. [Nate Tangsurat]
- See also: PSOBAT-1534
- Reformatted how security group is constructed. [Sijis Aviles]
 - Fix linting warning. [Sijis Aviles]
 - Use utils get_properties. [Sijis Aviles]
 - Merge branch ‘master’ of github.com:gogoair/foremast. [Sijis Aviles]
 - Docs: Improve help for *create-iam* [Nate Tangsurat]
 - Add ability to set security groups on instances. [Sijis Aviles]
 - Use get_properties util to get properties content. [Sijis Aviles]
 - Send complete file if no environment is specified. [Sijis Aviles]
 - Added build as an accepted environment. [Doug Campbell]
 - Update warning message to be more clear. [Sijis Aviles]

1.12.241 v2.2 (2016-05-12)

Features

- Add *destroy-elb* command. [Nate Tangsurat]

See also: PSOBAT-1534

Other

- V2.2. [Nate Tangsurat]
- Chore: Remove unused comment. [Nate Tangsurat]
See also: PSOBAT-1534
- Reformatted how security group is constructed. [Sijis Aviles]
- Fix linting warning. [Sijis Aviles]
- Use utils get_properties. [Sijis Aviles]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Sijis Aviles]
- Docs: Improve help for *create-iam* [Nate Tangsurat]
- Add ability to set security groups on instances. [Sijis Aviles]
- Use get_properties util to get properties content. [Sijis Aviles]
- Send complete file if no environment is specified. [Sijis Aviles]
- Added build as an accepted environment. [Doug Campbell]
- Update warning message to be more clear. [Sijis Aviles]

1.12.242 v2.1.0 (2017-05-31)

Fix

- Remove unused properties argument. [Nate Tangsurat]
See also: PSOBAT-1533
- Store HTTP verb as class attribute. [Nate Tangsurat]
See also: PSOBAT-1533

Features

- Add *destroy-sg* command. [Nate Tangsurat]
See also: PSOBAT-1533
- Add Gate API interface. [Nate Tangsurat]
See also: PSOBAT-1533

Other

- V2.1. [Nate Tangsurat]
- Docs: Expand on Gate() docstring. [Nate Tangsurat]
See also: PSOBAT-1533

1.12.243 v2.1 (2016-05-11)

Fix

- Remove unused properties argument. [Nate Tangsurat]

See also: PSOBAT-1533

- Store HTTP verb as class attribute. [Nate Tangsurat]

See also: PSOBAT-1533

Features

- Add *destroy-sg* command. [Nate Tangsurat]

See also: PSOBAT-1533

- Add Gate API interface. [Nate Tangsurat]

See also: PSOBAT-1533

Other

- V2.1. [Nate Tangsurat]

- Docs: Expand on Gate() docstring. [Nate Tangsurat]

See also: PSOBAT-1533

1.12.244 v2.0.0 (2017-05-31)

Fix

- Singular file name get_env_credentials.py. [Nate Tangsurat]

See also: PSOBAT-1544

- Do not attach 2 125GB volumes to each deploy. [Sijis Aviles]

The default setup for spinnaker is to attach 2 12GB volumes to each deployment. In our case, this is unnecessary and overkill.

Related spinnaker bug: <https://github.com/spinnaker/spinnaker/issues/721>

- Use Region for boto3 EC2 clients. [Nate Tangsurat]

- Stringify error to get message. [Nate Tangsurat]

See also: PSOBAT-1528

- Verbose DEBUG statements. [Nate Tangsurat]

See also: PSOBAT-1528

Features

- Generate ELB SSL certificate ARN. [Nate Tangsurat]

When a *certificate* has been specified for an ELB port, the name will be formatted into a fully qualified ARN with the Account number for the given Environment.

See also: PSOBAT-1544

- Add *destroy-s3* command. [Nate Tangsurat]

See also: PSOBAT-1532

Other

- V2.0. [Nate Tangsurat]

- Docs: Missing default from docstring. [Nate Tangsurat]

See also: PSOBAT-1544

- Merge branch ‘master’ of github.com:gogoair/foremast. [Sijis Aviles]

- Chore: TODO to retry DNS boto3 call. [Nate Tangsurat]

- Merge branch ‘master’ of github.com:gogoair/foremast. [Doug Campbell]

- Fixed bug with tagging group name, sorry for all the commits to master. [Doug Campbell]

- Testing bug with templates. [Doug Campbell]

- Better logging. [Doug Campbell]

- Fixed groupname tag for instances. [Doug Campbell]

- Merge branch ‘master’ of github.com:gogoair/foremast. [Doug Campbell]

- Merge branch ‘feature/security_group_cidr’ [Sijis Aviles]

- Add warning to user about older security groups. [Sijis Aviles]

- Docs: Fix example, protocol is TCP or UDP. [Nate Tangsurat]

See also: PSOBAT-1528

- Docs: Add Security Group example. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Avoid * imports and YAPF. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Format strings should reference index. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: YAPF. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Avoid variable reuse. [Nate Tangsurat]

Use more verbose variable name and avoid hidden issues when reusing variables.

See also: PSOBAT-1528

- Docs: Enrich docstrings. [Nate Tangsurat]
See also: PSOBAT-1528
- Add cidr rules to existing security group. [Sijis Aviles]
Spinnaker does not natively allow cidr to be specified, so this is using aws api directly to add that functionality.
- Fix pep8 warnings. [Sijis Aviles]
- Add util to get the id of a security group. [Sijis Aviles]
- Add generic security group exception. [Sijis Aviles]
- Add cidr functionality and validation. [Sijis Aviles]
- Fix logging. [Sijis Aviles]
- Added append for securitygroups.elb_extras. [Doug Campbell]
- Removed elb sg arg and added app sg. [Doug Campbell]

1.12.245 v2.0 (2016-05-11)

Fix

- Singular file name get_env_credentials.py. [Nate Tangsurat]
See also: PSOBAT-1544
- Do not attach 2 125GB volumes to each deploy. [Sijis Aviles]
The default setup for spinnaker is to attach 2 12GB volumes to each deployment. In our case, this is unnecessary and overkill.
Related spinnaker bug: <https://github.com/spinnaker/spinnaker/issues/721>
- Use Region for boto3 EC2 clients. [Nate Tangsurat]
- Stringify error to get message. [Nate Tangsurat]
See also: PSOBAT-1528
- Verbose DEBUG statements. [Nate Tangsurat]
See also: PSOBAT-1528

Features

- Generate ELB SSL certificate ARN. [Nate Tangsurat]
When a *certificate* has been specified for an ELB port, the name will be formatted into a fully qualified ARN with the Account number for the given Environment.
See also: PSOBAT-1544
- Add *destroy-s3* command. [Nate Tangsurat]
See also: PSOBAT-1532

Other

- V2.0. [Nate Tangsurat]
- Docs: Missing default from docstring. [Nate Tangsurat]

See also: PSOBAT-1544

- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Sijis Aviles]
- Chore: TODO to retry DNS boto3 call. [Nate Tangsurat]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Doug Campbell]
- Fixed bug with tagging group name, sorry for all the commits to master. [Doug Campbell]
- Testing bug with templates. [Doug Campbell]
- Better logging. [Doug Campbell]
- Fixed groupname tag for instances. [Doug Campbell]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Doug Campbell]
- Merge branch ‘feature/security_group_cidr’ [Sijis Aviles]
- Add warning to user about older security groups. [Sijis Aviles]
- Docs: Fix example, protocol is TCP or UDP. [Nate Tangsurat]

See also: PSOBAT-1528

- Docs: Add Security Group example. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Avoid * imports and YAPF. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Format strings should reference index. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: YAPF. [Nate Tangsurat]

See also: PSOBAT-1528

- Style: Avoid variable reuse. [Nate Tangsurat]

Use more verbose variable name and avoid hidden issues when reusing variables.

See also: PSOBAT-1528

- Docs: Enrich docstrings. [Nate Tangsurat]

See also: PSOBAT-1528

- Add cidr rules to existing security group. [Sijis Aviles]

Spinnaker does not natively allow cidr to be specified, so this is using aws api directly to add that functionality.

- Fix pep8 warnings. [Sijis Aviles]

- Add util to get the id of a security group. [Sijis Aviles]

- Add generic security group exception. [Sijis Aviles]

- Add cidr functionality and validation. [Sijis Aviles]

- Fix logging. [Sijis Aviles]

- Added append for securitygroups.elb_extras. [Doug Campbell]
- Removed elb sg arg and added app sg. [Doug Campbell]

1.12.246 v1.8.0 (2017-05-31)

Fix

- Forgot __init__ [Nate Tangsurat]

See also: PSOBAT-1531

Other

- V1.8. [Nate Tangsurat]

1.12.247 v1.8 (2016-05-10)

Fix

- Forgot __init__ [Nate Tangsurat]

See also: PSOBAT-1531

Other

- V1.8. [Nate Tangsurat]

1.12.248 v1.7.0 (2017-05-31)

Fix

- Remove intermediary log message prefix. [Nate Tangsurat]

See also: PSOBAT-1531

- WARNING when no environment settings found. [Nate Tangsurat]

See also: PSOBAT-1482

- Handle no “services” in pipeline.json. [Nate Tangsurat]

See also: PSOBAT-1482

- Restrict S3 access to read only. [Nate Tangsurat]

See also: PSOBAT-1482

- Remove unused import. [Nate Tangsurat]

- Remove one more argument. [Nate Tangsurat]

See also: PSOBAT-1444

- Remove more deprecated arguments. [Nate Tangsurat]

See also: PSOBAT-1444

- Remove extra *create-elb* argument. [Nate Tangsurat]
See also: PSOBAT-1444
- Add default null certificate to old configs. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove extra JSON dumps/loads. [Nate Tangsurat]
See also: PSOBAT-1444
- INFO each ELB Listener found. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1444

Features

- Add *destroy-iam* command. [Nate Tangsurat]
See also: PSOBAT-1531
- Delete IAM Resources function. [Nate Tangsurat]
See also: PSOBAT-1531
- Add inline IAM Policy to Role. [Nate Tangsurat]
Based on *pipeline.json services* key, an inline Policy will be attached to the IAM Role.
See also: PSOBAT-1482
- Construct IAM Policies. [Nate Tangsurat]
See also: PSOBAT-1482

Other

- V1.7. [Nate Tangsurat]
- Chore: Remove configurations in favour of consts. [Nate Tangsurat]
- Refactor: Unpack to keywords directly. [Nate Tangsurat]
- Style: Sort add_*() functions. [Nate Tangsurat]
- Refactor: Use add_gitlab_token() [Nate Tangsurat]
- Refactor: Use add_properties() [Nate Tangsurat]
- Refactor: Use add_region() [Nate Tangsurat]
- Refactor: Use add_env() [Nate Tangsurat]
- Refactor: Switch to add_app() [Nate Tangsurat]
- Refactor: Pull out available environments. [Nate Tangsurat]
See also: PSOBAT-1531
- Refactor: Pull out arguments for reuse. [Nate Tangsurat]
See also: PSOBAT-1531

- Refactor: Pull out resource_action() for reuse. [Nate Tangsurat]

See also: PSOBAT-1531

- Tests: construct_policy() can return None. [Nate Tangsurat]

See also: PSOBAT-1482

- Tests: Test IAM Policy with no “services” [Nate Tangsurat]

See also: PSOBAT-1482

- Chore: Add TODO for more IAM Policy testing. [Nate Tangsurat]

See also: PSOBAT-1482

- Docs: Fix example of services for IAM Policies. [Nate Tangsurat]

See also: PSOBAT-1482

- Tests: Fix S3 Policy test. [Nate Tangsurat]

See also: PSOBAT-1482

- Merge branch ‘features/security_group’ [Sijis Aviles]

- Ensure the advanced style is a list. [Sijis Aviles]

This handles situations where an application has multiple ports defined but they are different protocols

- Add apps_all to default security group. [Sijis Aviles]

- Handle older published securitygroup format. [Sijis Aviles]

We essentially have two security group formats: - simple: just a list of ports

```
{ 'ingress': { 'app': [1,2,3,4] } }
```

– advanced: ports ranges { ‘ingress’: { ‘app’: { ‘start_port’: 1, ‘end_port’: 10, ‘protocol’: ‘udp’ } } }

- Split ports into separate distinct keys. [Sijis Aviles]

- Use given description for security group. [Sijis Aviles]

- Update security group based on property file values. [Sijis Aviles]

- Add a properties file parameter for securitygroup. [Sijis Aviles]

- Add get_properties util method. [Sijis Aviles]

- Forgot a comma. [Doug Campbell]

- More typo fixes. [Doug Campbell]

- Fixed json typo. [Doug Campbell]

- Added notification to templates. [Doug Campbell]

- Set healthcheck grace period to 5 minutes. [Sijis Aviles]

- Removed bypassing elb health check amazon stuff. [Doug Campbell]

- Updated jenkins job name for audit-approval. [Doug Campbell]

- Fixed audit-approval json to fail properly and added a stage comment. [Doug Campbell]

- Changed log-deployment to audit approval and fixed renumerate to work with it. [Doug Campbell]

- Tests: Make sure ‘ports’ key overrides. [Nate Tangsurat]

See also: PSOBAT-1444

- Tests: Remove unneeded key by key comparison. [Nate Tangsurat]
See also: PSOBAT-1444
- Tests: Add test for format_listeners() [Nate Tangsurat]
See also: PSOBAT-1444
- Tests: Update splay_health() test to new module. [Nate Tangsurat]
- Refactor: Extract reused ELB settings lookup. [Nate Tangsurat]
See also: PSOBAT-1444
- Docs: Update format_listeners() docstring example. [Nate Tangsurat]
See also: PSOBAT-1444
- Refactor: Separate out Listener reading logic. [Nate Tangsurat]
See also: PSOBAT-1444

1.12.249 v1.7 (2016-05-10)

- V1.7. [Nate Tangsurat]

1.12.250 v1.6 (2016-05-10)

- Chore: Remove configurations in favour of consts. [Nate Tangsurat]
- Refactor: Unpack to keywords directly. [Nate Tangsurat]

1.12.251 v1.5 (2016-05-10)

- Style: Sort add_*() functions. [Nate Tangsurat]
- Refactor: Use add_gitlab_token() [Nate Tangsurat]
- Refactor: Use add_properties() [Nate Tangsurat]
- Refactor: Use add_region() [Nate Tangsurat]
- Refactor: Use add_env() [Nate Tangsurat]
- Refactor: Switch to add_app() [Nate Tangsurat]

1.12.252 v1.4 (2016-05-10)

Fix

- Remove intermediary log message prefix. [Nate Tangsurat]
See also: PSOBAT-1531

Features

- Add *destroy-iam* command. [Nate Tangsurat]
See also: PSOBAT-1531
- Delete IAM Resources function. [Nate Tangsurat]
See also: PSOBAT-1531

Other

- Refactor: Pull out available environments. [Nate Tangsurat]
See also: PSOBAT-1531
- Refactor: Pull out arguments for reuse. [Nate Tangsurat]
See also: PSOBAT-1531
- Refactor: Pull out resource_action() for reuse. [Nate Tangsurat]
See also: PSOBAT-1531

1.12.253 v1.3 (2016-05-09)

Fix

- WARNING when no environment settings found. [Nate Tangsurat]
See also: PSOBAT-1482
- Handle no “services” in pipeline.json. [Nate Tangsurat]
See also: PSOBAT-1482
- Restrict S3 access to read only. [Nate Tangsurat]
See also: PSOBAT-1482
- Remove unused import. [Nate Tangsurat]

Features

- Add inline IAM Policy to Role. [Nate Tangsurat]
Based on *pipeline.json services* key, an inline Policy will be attached to the IAM Role.
See also: PSOBAT-1482
- Construct IAM Policies. [Nate Tangsurat]
See also: PSOBAT-1482

Other

- Tests: construct_policy() can return None. [Nate Tangsurat]
See also: PSOBAT-1482

- Tests: Test IAM Policy with no “services” [Nate Tangsurat]
See also: PSOBAT-1482
- Chore: Add TODO for more IAM Policy testing. [Nate Tangsurat]
See also: PSOBAT-1482
- Docs: Fix example of services for IAM Policies. [Nate Tangsurat]
See also: PSOBAT-1482
- Tests: Fix S3 Policy test. [Nate Tangsurat]
See also: PSOBAT-1482

1.12.254 v1.2 (2016-05-09)

Fix

- Remove one more argument. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove more deprecated arguments. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove extra *create-elb* argument. [Nate Tangsurat]
See also: PSOBAT-1444
- Add default null certificate to old configs. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove extra JSON dumps/loads. [Nate Tangsurat]
See also: PSOBAT-1444
- INFO each ELB Listener found. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1444
- Add default ELB protocols at template level. [Nate Tangsurat]
See also: PSOBAT-1444
- Use Health Check target from properties. [Nate Tangsurat]
See also: PSOBAT-1444
- Use subnet purose from properties. [Nate Tangsurat]
See also: PSOBAT-1444
- Remove unused arguments. [Nate Tangsurat]
See also: PSOBAT-1444
- Use properties to fill in ELB Listeners. [Nate Tangsurat]
See also: PSOBAT-1444

- Use *create-configs* output for properties. [Nate Tangsurat]

See also: PSOBAT-1444

- Use correct generic custom Base OS. [Nate Tangsurat]

See also: PSOBAT-1468

- Use Base OS “other” [Nate Tangsurat]

See also: PSOBAT-1468

Other

- Merge branch ‘features/security_group’ [Sijis Aviles]

- Ensure the advanced style is a list. [Sijis Aviles]

This handles situations where an application has multiple ports defined but they are different protocols

- Add apps_all to default security group. [Sijis Aviles]

- Handle older published securitygroup format. [Sijis Aviles]

We essentially have two security group formats: - simple: just a list of ports

```
{ 'ingress': { 'app': [1,2,3,4] } }
```

– advanced: ports ranges { ‘ingress’: { ‘app’: { ‘start_port’: 1, ‘end_port’: 10, ‘protocol’: ‘udp’ } } }

- Split ports into separate distinct keys. [Sijis Aviles]

- Use given description for security group. [Sijis Aviles]

- Update security group based on property file values. [Sijis Aviles]

- Add a properties file parameter for securitygroup. [Sijis Aviles]

- Add get_properties util method. [Sijis Aviles]

- Forgot a comma. [Doug Campbell]

- More typo fixes. [Doug Campbell]

- Fixed json typo. [Doug Campbell]

- Added notification to templates. [Doug Campbell]

- Set healthcheck grace period to 5 minutes. [Sijis Aviles]

- Removed bypassing elb health check amazon stuff. [Doug Campbell]

- Updated jenkins job name for audit-approval. [Doug Campbell]

- Fixed audit-approval json to fail properly and added a stage comment. [Doug Campbell]

- Changed log-deployment to audit approval and fixed renumerate to work with it. [Doug Campbell]

- Tests: Make sure ‘ports’ key overrides. [Nate Tangsurat]

See also: PSOBAT-1444

- Tests: Remove unneeded key by key comparison. [Nate Tangsurat]

See also: PSOBAT-1444

- Tests: Add test for format_listeners() [Nate Tangsurat]

See also: PSOBAT-1444

- Tests: Update splay_health() test to new module. [Nate Tangsurat]
See also: PSOBAT-1444
- Refactor: Extract reused ELB settings lookup. [Nate Tangsurat]
See also: PSOBAT-1444
- Docs: Update format_listeners() docstring example. [Nate Tangsurat]
See also: PSOBAT-1444
- Refactor: Separate out Listener reading logic. [Nate Tangsurat]
See also: PSOBAT-1444
- Style: Prettify and sort. [Nate Tangsurat]
See also: PSOBAT-1444
- Refactor: Move splay_health() into file. [Nate Tangsurat]
See also: PSOBAT-1444

1.12.255 v1.1.0 (2017-05-31)

Fix

- Remove “qe” contents. [Nate Tangsurat]
See also: PSOBAT-1448
- Send Python dict to QE. [Nate Tangsurat]
See also: PSOBAT-1448
- Remove duplicate template. [Nate Tangsurat]
See also: PSOBAT-1468
- Expand user path. [Nate Tangsurat]
See also: PSOBAT-1468
- Use default of \$HOME. [Nate Tangsurat]
See also: PSOBAT-1468
- INFO message env and region should be flipped. [Nate Tangsurat]
See also: PSOBAT-1468
- Use safer dict.get() [Nate Tangsurat]
See also: PSOBAT-1468
- Use AMI based on name and Region. [Nate Tangsurat]
See also: PSOBAT-1468
- Default to tomcat8. [Nate Tangsurat]
See also: PSOBAT-1468
- Use raw.properties.json for Base OS. [Nate Tangsurat]
See also: PSOBAT-1468

- Strip out any quotes for –base. [Nate Tangsurat]

See also: PSOBAT-1468

- Add “base” option to pipeline.json. [Nate Tangsurat]

Base AMI will default to “tomcat” if not specified. Also added default deployment environments.

See also: PSOBAT-1468

- Remove trailing comma. [Nate Tangsurat]

Trailing commas turn normal assignments into tuple assignments, very hidden and confusing.

See also: PSOBAT-1197

- Entry point command does not need *python* [Nate Tangsurat]

See also: PSOBAT-1197

- Pip install foremast. [Nate Tangsurat]

See also: PSOBAT-1197

- Use *create-app* command for *pipes-app* [Nate Tangsurat]

See also: PSOBAT-1197

- One reference to pipeline.json environments. [Nate Tangsurat]

See also: PSOBAT-1197

- INFO message refers to blocks. [Nate Tangsurat]

See also: PSOBAT-1197

- Prettify found envs and regions for Pipelines. [Nate Tangsurat]

See also: PSOBAT-1197

- Put failure message in exception. [Nate Tangsurat]

See also: PSOBAT-1197

- Remove unused import. [Nate Tangsurat]

See also: PSOBAT-1197

- Region parameter should not be modified. [Nate Tangsurat]

See also: PSOBAT-1197

- Slightly cleaner environment list. [Nate Tangsurat]

See also: PSOBAT-1197

- Use more concise error message for VPCs. [Nate Tangsurat]

See also: PSOBAT-1197

- Use generated details in clean_pipelines() [Nate Tangsurat]

See also: PSOBAT-1197

- Use generated Instance Profile. [Nate Tangsurat]

Default to the generated Instance Profile name instead of “app_default_profile”.

- Remove pci and sox. [Nate Tangsurat]

Force usages of bad names to fail now while we’re transitioning.

- Format Health Check path with slashes. [Nate Tangsurat]
See also: PSOBAT-1113
- INFO Health Check properties. [Nate Tangsurat]
See also: PSOBAT-1113
- Make AccessDenied errors more evident. [Nate Tangsurat]
See also: PSOBAT-1113
- Use extended sequence unpacking for clarity. [Nate Tangsurat]
See also: PSOBAT-1359
- Remove alerting for non-failures. [Nate Tangsurat]
See also: PSOBAT-1359
- Simplify Stage reference IDs. [Nate Tangsurat]
See also: PSOBAT-1359
- Hard code Git taggers. [Nate Tangsurat]
See also: PSOBAT-1359
- DEBUG Regions dict. [Nate Tangsurat]
See also: PSOBAT-1359

Features

- Add Base AMI option for Bake Stage. [Nate Tangsurat]
See also: PSOBAT-1468

Other

- V1.1. [Nate Tangsurat]
- Added passing of QE JSON to downstream job. [Joel Vasallo]
- Chore: Remove unused templates. [Nate Tangsurat]
See also: PSOBAT-1468
- Tests: More lenient test. [Nate Tangsurat]
AMI IDs will change in the future. All we care about is that we get an AMI ID.
See also: PSOBAT-1468
- Tests: ami_lookup() testing. [Nate Tangsurat]
See also: PSOBAT-1468
- Update jenkins jobs to use simple bash script. [Sijis Aviles]
- Merge branch ‘fix/clean_pipes’ [Sijis Aviles]
- Added minor TODO. [Sijis Aviles]
- Docs: Update running instructions. [Nate Tangsurat]
See also: PSOBAT-1197

- Docs: Can run code from root of repository. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Remove unused import. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Remove extra INFO. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Reduce number of calls to get_subnets() [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Remove some extra logging. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: No longer need regions in blocks. [Nate Tangsurat]
The `data.app.regions` template variable exists in the Find Images Stage, which is no longer used with the decompiled Pipeline blocks.
See also: PSOBAT-1197
- Style: Use more descriptive variable names. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Extract construct_pipeline_block() [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Simplify blocks and extract functions. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Convert get_settings() to staticmethod. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: get_pipe_id() -> get_pipeline_id() [Nate Tangsurat]
See also: PSOBAT-1197
- Chore: Remove unneeded comments. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: get_all_pipelines(), clean_pipelines() [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Run create-pipeline only once at end. [Nate Tangsurat]
See also: PSOBAT-1197
- Style: Retab. [Nate Tangsurat]
See also: PSOBAT-1197
- Updated runway/dsl for app job. [Doug Campbell]
- Revert “cleaned up jenkins jobs” [Doug Campbell]
This reverts commit c9a044c6970c652c5ec5b1bb33805847996d3b38.
- Cleaned up jenkins jobs. [Doug Campbell]

- Updated combine loop to account for deployment logs. [Doug Campbell]
- Updated template for processor jenkins job. [Doug Campbell]
- Fixed directory naming. [Doug Campbell]
- Added name tag to templates. [Doug Campbell]
- Merge branch ‘master’ of github.com:gogoair/foremast. [Doug Campbell]
- Fixed key naming for prods and prodp. [Doug Campbell]
- Changed prodsox/prodpci to props/prodp. [Doug Campbell]
- Fixed s3 prodsox/prodpci. [Doug Campbell]
- Merge branch ‘master’ of github.com:gogoair/foremast. [Doug Campbell]
- Changed pci/sox to prodpci/prodsox. [Doug Campbell]
- More typo fixes in tempaltes. [Doug Campbell]
- Fixed typo. [Doug Campbell]
- Fixed issues in templates again. [Doug Campbell]
- Fixed elb naming int templates. [Doug Campbell]
- Fixed min instance issue. [Doug Campbell]
- Changed pipeline name. [Doug Campbell]
- Docs: Add refId example for Stages. [Nate Tangsurat]

See also: PSOBAT-1359

1.12.256 v1.1 (2016-05-03)

Fix

- Remove “qe” contents. [Nate Tangsurat]
See also: PSOBAT-1448
- Send Python dict to QE. [Nate Tangsurat]
See also: PSOBAT-1448
- Remove duplicate template. [Nate Tangsurat]
See also: PSOBAT-1468
- Expand user path. [Nate Tangsurat]
See also: PSOBAT-1468
- Use default of \$HOME. [Nate Tangsurat]
See also: PSOBAT-1468
- INFO message env and region should be flipped. [Nate Tangsurat]
See also: PSOBAT-1468
- Use safer dict.get() [Nate Tangsurat]
See also: PSOBAT-1468

- Use AMI based on name and Region. [Nate Tangsurat]

See also: PSOBAT-1468

- Default to tomcat8. [Nate Tangsurat]

See also: PSOBAT-1468

- Use raw.properties.json for Base OS. [Nate Tangsurat]

See also: PSOBAT-1468

- Strip out any quotes for –base. [Nate Tangsurat]

See also: PSOBAT-1468

- Add “base” option to pipeline.json. [Nate Tangsurat]

Base AMI will default to “tomcat” if not specified. Also added default deployment environments.

See also: PSOBAT-1468

- Remove trailing comma. [Nate Tangsurat]

Trailing commas turn normal assignments into tuple assignments, very hidden and confusing.

See also: PSOBAT-1197

- Entry point command does not need *python* [Nate Tangsurat]

See also: PSOBAT-1197

- Pip install foremast. [Nate Tangsurat]

See also: PSOBAT-1197

- Use *create-app* command for *pipes-app* [Nate Tangsurat]

See also: PSOBAT-1197

- One reference to pipeline.json environments. [Nate Tangsurat]

See also: PSOBAT-1197

- INFO message refers to blocks. [Nate Tangsurat]

See also: PSOBAT-1197

- Prettify found envs and regions for Pipelines. [Nate Tangsurat]

See also: PSOBAT-1197

- Put failure message in exception. [Nate Tangsurat]

See also: PSOBAT-1197

- Remove unused import. [Nate Tangsurat]

See also: PSOBAT-1197

- Region parameter should not be modified. [Nate Tangsurat]

See also: PSOBAT-1197

- Slightly cleaner environment list. [Nate Tangsurat]

See also: PSOBAT-1197

- Use more concise error message for VPCs. [Nate Tangsurat]

See also: PSOBAT-1197

- Use generated details in clean_pipelines() [Nate Tangsurat]

See also: PSOBAT-1197

Features

- Add Base AMI option for Bake Stage. [Nate Tangsurat]

See also: PSOBAT-1468

Other

- V1.1. [Nate Tangsurat]

- Added passing of QE JSON to downstream job. [Joel Vasallo]

- Chore: Remove unused templates. [Nate Tangsurat]

See also: PSOBAT-1468

- Tests: More lenient test. [Nate Tangsurat]

AMI IDs will change in the future. All we care about is that we get an AMI ID.

See also: PSOBAT-1468

- Tests: ami_lookup() testing. [Nate Tangsurat]

See also: PSOBAT-1468

- Update jenkins jobs to use simple bash script. [Sijis Aviles]

- Merge branch ‘fix/clean_pipes’ [Sijis Aviles]

- Added minor TODO. [Sijis Aviles]

- Docs: Update running instructions. [Nate Tangsurat]

See also: PSOBAT-1197

- Docs: Can run code from root of repository. [Nate Tangsurat]

See also: PSOBAT-1197

- Refactor: Remove unused import. [Nate Tangsurat]

See also: PSOBAT-1197

- Refactor: Remove extra INFO. [Nate Tangsurat]

See also: PSOBAT-1197

- Refactor: Reduce number of calls to get_subnets() [Nate Tangsurat]

See also: PSOBAT-1197

- Refactor: Remove some extra logging. [Nate Tangsurat]

See also: PSOBAT-1197

- Refactor: No longer need regions in blocks. [Nate Tangsurat]

The `data.app.regions` template variable exists in the Find Images Stage, which is no longer used with the de-compiled Pipeline blocks.

See also: PSOBAT-1197

- Style: Use more descriptive variable names. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Extract construct_pipeline_block() [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Simplify blocks and extract functions. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Convert get_settings() to staticmethod. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: get_pipe_id() -> get_pipeline_id() [Nate Tangsurat]
See also: PSOBAT-1197
- Chore: Remove unneeded comments. [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: get_all_pipelines(), clean_pipelines() [Nate Tangsurat]
See also: PSOBAT-1197
- Refactor: Run create-pipeline only once at end. [Nate Tangsurat]
See also: PSOBAT-1197
- Style: Retab. [Nate Tangsurat]
See also: PSOBAT-1197

1.12.257 v0.0 (2016-04-28)

Fix

- Use generated Instance Profile. [Nate Tangsurat]
Default to the generated Instance Profile name instead of “app_default_profile”.
- Remove pci and sox. [Nate Tangsurat]
Force usages of bad names to fail now while we’re transitioning.
- Format Health Check path with slashes. [Nate Tangsurat]
See also: PSOBAT-1113
- INFO Health Check properties. [Nate Tangsurat]
See also: PSOBAT-1113
- Make AccessDenied errors more evident. [Nate Tangsurat]
See also: PSOBAT-1113
- Use extended sequence unpacking for clarity. [Nate Tangsurat]
See also: PSOBAT-1359
- Remove alerting for non-failures. [Nate Tangsurat]
See also: PSOBAT-1359

- Simplify Stage reference IDs. [Nate Tangsurat]
See also: PSOBAT-1359
- Hard code Git taggers. [Nate Tangsurat]
See also: PSOBAT-1359
- DEBUG Regions dict. [Nate Tangsurat]
See also: PSOBAT-1359
- Stage names corrected. [Nate Tangsurat]
See also: PSOBAT-1359
- Lower case Pipeline names. [Nate Tangsurat]
See also: PSOBAT-1372
- Minimize Pipeline notifications. [Nate Tangsurat]
See also: PSOBAT-1374
- Convert bytes to string. [Nate Tangsurat]
See also: PSOBAT-1374
- Include User Data. [Nate Tangsurat]
See also: PSOBAT-1374
- Add ELB and SG name to templates. [Nate Tangsurat]
- Update Manual Judgement comment. [Nate Tangsurat]
See also: PSOBAT-1361
- Use InstanceProfileName value to remove. [Nate Tangsurat]
- Start using master Branch again. [Nate Tangsurat]
See also: PSOBAT-1250

Features

- Add Git Tagger stages. [Nate Tangsurat]
See also: PSOBAT-1359

Other

- Updated runway/dsl for app job. [Doug Campbell]
- Revert “cleaned up jenkins jobs” [Doug Campbell]
This reverts commit c9a044c6970c652c5ec5b1bb33805847996d3b38.
- Cleaned up jenkins jobs. [Doug Campbell]
- Updated combine loop to account for deployment logs. [Doug Campbell]
- Updated template for processor jenkins job. [Doug Campbell]
- Fixed directory naming. [Doug Campbell]
- Added name tag to templates. [Doug Campbell]

- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Doug Campbell]
- Fixed key naming for prods and prodp. [Doug Campbell]
- Changed prodsox/prodpci to props/prodp. [Doug Campbell]
- Fixed s3 prodsox/prodpci. [Doug Campbell]
- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Doug Campbell]
- Changed pci/sox to prodpci/prodsox. [Doug Campbell]
- More typo fixes in tempaltes. [Doug Campbell]
- Fixed typo. [Doug Campbell]
- Fixed issues in templates again. [Doug Campbell]
- Fixed elb naming int templates. [Doug Campbell]
- Fixed min instance issue. [Doug Campbell]
- Changed pipeline name. [Doug Campbell]
- Docs: Add refId example for Stages. [Nate Tangsurat]

See also: PSOBAT-1359

- Style: Fix up whitespace and docstring. [Nate Tangsurat]

See also: PSOBAT-1359

- Fixed big with region looping. [Doug Campbell]
- Added region support. [Doug Campbell]
- Added PCI template. [Doug Campbell]
- Fixed templating, added sox specific template. [Doug Campbell]
- Updated logic to combine pipelines. [Doug Campbell]
- Chore: Add TODO to consider using defaults always. [Nate Tangsurat]

sbasgall posited that when an *application-{branch}-{env}.json* file is missing, should we default to the templated configuration? This could introduce the issue of people not being aware that they are using the defaults instead of needing to create the JSON file, but may be worth it.

- Docs: Spruce up docstrings. [Nate Tangsurat]

See also: PSOBAT-1374

- Docs: Finished Manual Judgement. [Nate Tangsurat]

See also: PSOBAT-1361

- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Sijis Aviles]
- Add all accounts when new application is creation. [Sijis Aviles]
- Docs: Add more TODOs. [Nate Tangsurat]

See also: PSOBAT-1250

- Merge branch ‘master’ of git.gogoair.com:spinnaker/pipes. [Sijis Aviles]
- Docs: Update with initial migration. [Nate Tangsurat]

See also: PSOBAT-1250

- Fix desired capacity for trigger template. [Sijis Aviles]

- Desired capacity should use min_inst defined. [Sijis Aviles]
- Added sox and pci to args. [Doug Campbell]

CHAPTER 2

Spinnaker Foremast

Foremast is a Spinnaker pipeline and infrastructure configuration and templating tool. Just create a couple JSON configuration files and then manually creating Spinnaker pipelines becomes a thing of the past.



2.1 Why Foremast?

- No manual creation of pipelines in the Spinnaker UI
- Reproducible and versioned Spinnaker pipelines
- Standardized pipelines with flexibility for application specific needs

With Foremast, Developers create a couple simple JSON configs per application. These configs provide details on the pipeline and infrastructure specific to the application's needs. Foremast takes those configs, renders some Jinja2 templates, and then acts as a client for the Spinnaker Gate API. Foremast comes with generic templates for creating a simple pipeline but it can also point to external templates for custom pipelines that fit any workflow.

2.2 Foremast Features

- Dynamically generate Spinnaker pipelines based on JSON configs
- Customizable pipelines through external Jinja2 Templates, see [Foremast templates](#) for examples
- Dynamically generate AWS infrastructure based on pipeline configs
- Set up resources not defined in Spinnaker, such as S3 buckets and IAM roles
- Support for AWS Lambda pipelines

2.3 Getting Started

Take a look at [quick start guide](#) for a quick introduction on how to use Foremast.

We also have a blog post to help you get started: [Automate Spinnaker Pipeline Creation](#)

2.3.1 Documentation

All the documentation can be viewed on [Read the Docs](#). You can find all configuration options, code information, and better examples there.

2.3.2 Development

See the [contribution guide](#) for information on code style, contributing, and testing.

2.3.3 Getting Help

For questions, support, or friendly conversation you can find us on [Gitter](#).

2.4 More Details

2.4.1 Installing

Installing the package will provide CLI commands for convenience.

```
virtualenv -p python3 venv
source venv/bin/activate
pip install foremast
```

2.4.2 Entry Points

Foremast has a few easy to use CLI endpoints.

- `foremast-pipeline` - Creates an application and pipeline Spinnaker
- `foremast-infrastructure` - Sets up AWS infrastructure like s3, iam, elb, and security groups
- `foremast-pipeline-onetime` - Generates a pipeline for deploying to one specific account

- `foremast-scaling-policy` - Creates and attaches a scaling policy to an application server group.
- `foremast-pipeline-rebuild` - rebuild pipelines after changes have been made

You can run any of these entries points from the command line. They rely on environment variables and are ideal for running in a Jenkins job

```
PROJECT=forrest GIT_REPO=core RUNWAY_DIR=path/to/pipeline_configs foremast-pipeline
```

2.4.3 Foremast Configuration

A file at `{pwd}/.foremast/foremast.cfg`, `~/.foremast/foremast.cfg`, or `/etc/foremast/foremast.cfg` needs to exist in order to run foremast.

```
[base]
domain = example.com
envs = dev,stage,prod
regions = us-east-1
gate_api_url = http://gate.example.com:8084
```

2.4.4 Runway Configuration Files

To begin using Foremast, you must have a few JSON configuration files defined for each application

`pipeline.json`

This file will be needed for each application. Foremast has a lot of defaults in place for `pipeline.json`, take a look at the [pipeline.json](#) docs for all options.

Minimum

```
{
    "deployment": "spinnaker"
}
```

Example Deployment Environments Override

Custom deployment environment order and selection can be provided in the `env` key. When missing, the default provided is `{"env": ["stage", "prod"]}`. Here, the order matters and Pipeline will be generated in the given order.

```
{
    "deployment": "spinnaker",
    "env": [
        "prod"
    ]
}
```

`application-master-{env}.json`

Each deployment environment specified in the `pipeline.json` file will need an accompanying `application-master-{env}.json` file in the same directory.

The ‘application-master-{env}’ files have a lot of exposed values with sane defaults. Please take a look at the application.json docs for all options.

application-master-{env}.json example

```
{  
    "security_group": {  
        "description": "something useful",  
        "elb_extras": ["sg_offices"],  
        "ingress": {},  
        "egress": "0.0.0.0/0"  
    },  
    "app": {  
        "instance_type": "t2.small",  
        "app_description": "Edge Forrest Demo application",  
        "instance_profile": "forrest_edge_profile"  
    },  
    "elb": {  
        "subnet_purpose": "internal",  
        "target": "TCP:8080",  
        "ports": [  
            {"loadbalancer": "HTTP:80", "instance": "HTTP:8080"}  
        ]  
    },  
    "asg": {  
        "subnet_purpose": "internal",  
        "min_inst": 1,  
        "max_inst": 1,  
        "scaling_policy": {  
            "metric": "CPUUtilization",  
            "threshold": 90,  
            "period_minutes": 10,  
            "statistic": "Average"  
        }  
    },  
    "regions": ["us-east-1"],  
    "dns" : {  
        "ttl": 120  
    }  
}
```

CHAPTER 3

Indices and tables

- genindex
- modindex
- search