
policy*entry*

Feb 02, 2020

1	Overview	3
2	Comparison to other tools	9
3	Installation	13
4	Initialization	15
5	Writing IAM Policies	17
6	Downloading Policies	23
7	Analyzing Policies	25
8	Querying the Policy Database	31
9	Docker	35
10	Command cheat sheet	37
11	Terraform Demo	41
12	Terraform Modules	43
13	Contributing	45
14	Library Usage	53
15	Module Reference	55
16	Examples	67
17	Implementation Strategy	75
18	IAM Policies	77
19	Minimization	81
20	Indices and tables	83

Python Module Index	85
Index	87

Policy Sentry is an AWS IAM Least Privilege Policy Generator, auditor, and analysis database. It compiles database tables based on the AWS IAM Documentation on [Actions](#), [Resources](#), and [Condition Keys](#) and leverages that data to create least-privilege IAM policies.

Organizations can use Policy Sentry to:

- **Limit the blast radius in the event of a breach:** If an attacker gains access to user credentials or Instance Profile credentials, access levels and resource access should be limited to the least amount needed to function. This can help avoid situations such as the Capital One breach, where after an SSRF attack, data was accessible from the compromised instance because the role allowed access to all S3 buckets in the account. In this case, Policy Sentry would only allow the role access to the buckets necessary to perform its duties.
- **Scale creation of secure IAM Policies:** Rather than dedicating specialized and talented human resources to manual IAM reviews and creating IAM policies by hand, organizations can leverage Policy Sentry to write the policies for them in an automated fashion.

Policy Sentry's policy writing templates are expressed in YAML and include the following:

- Name and Justification for why the privileges are needed
- CRUD levels (Read/Write/List/Tagging/Permissions management)
- Amazon Resource Names (ARNs), so the resulting policy only points to specific resources and does not grant access to * resources.

Policy Sentry can also be used to:

- [Audit IAM Policies](#) based on access levels
- [Query the IAM database](#) to reduce manual search time
- [Download live policies](#) from an AWS account auditing purposes
- [Generate IAM Policies based on Terraform output](#)
- [Write least-privilege IAM Policies](#) based on a list of IAM actions (or CRUD levels)

Navigate below to get started with Policy Sentry!

Policy Sentry is an IAM Least Privilege Policy Generator, auditor, and analysis database.

1.1 Motivation

Writing security-conscious IAM Policies by hand can be very tedious and inefficient. Many Infrastructure as Code developers have experienced something like this:

- Determined to make your best effort to give users and roles the least amount of privilege you need to perform your duties, you spend way too much time combing through the AWS IAM Documentation on [Actions](#), [Resources](#), and [Condition Keys for AWS Services](#).
- Your team lead encourages you to build security into your IAM Policies for product quality, but eventually you get frustrated due to project deadlines.
- You don't have an embedded security person on your team who can write those IAM policies for you, and there's no automated tool that will automagically sense the AWS API calls that you perform and then write them for you in a least-privilege manner.
- After fantasizing about that level of automation, you realize that writing least privilege IAM Policies, seemingly out of charity, will jeopardize your ability to finish your code in time to meet project deadlines.
- You use Managed Policies (because hey, why not) or you eyeball the names of the API calls and use wildcards instead so you can move on with your life.

Such a process is not ideal for security or for Infrastructure as Code developers. We need to make it easier to write IAM Policies securely and abstract the complexity of writing least-privilege IAM policies. That's why I made this tool.

1.1.1 Authoring Secure IAM Policies

`policy_sentry`'s flagship feature is that it can create IAM policies based on resource ARNs and access levels. Our CRUD functionality takes the opinionated approach that IAC developers shouldn't have to understand the complexities of AWS IAM - we should abstract the complexity for them. In fact, developers should just be able to say...

- “I need Read/Write/List access to `arn:aws:s3:::example-org-sbx-vmimport`”
- “I need Permissions Management access to `arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret`”
- “I need Tagging access to `arn:aws:ssm:us-east-1:123456789012:parameter/test`”

...and our automation should create policies that correspond to those access levels.

How do we accomplish this? Well, `policy_sentry` leverages the AWS documentation on [Actions, Resources, and Condition Keys](#) documentation to look up the actions, access levels, and resource types, and generates policies according to the ARNs and access levels. Consider the table snippet below:

Actions	Access Level	Resource Types
<code>ssm:GetParameter</code>	Read	parameter
<code>ssm:DescribeParameters</code>	List	parameter
<code>ssm:PutParameter</code>	Write	parameter
<code>secretsmanager:PutResourcePolicy</code>	Permissions management	secret
<code>secretsmanager:TagResource</code>	Tagging	secret

Policy Sentry aggregates all of that documentation into a single database and uses that database to generate policies according to actions, resources, and access levels. To generate a policy according to resources and access levels, start by creating a template with this command so you can just fill out the ARNs:

```
policy_sentry create-template --name myRole --output-file crud.yml --template-type_
↪crud
```

It will generate a file like this:

```
policy_with_crud_levels:
- name: myRole
  description: '' # Insert description
  arn: '' # Insert the ARN of the role that will use this
  read:
    - '' # Insert ARNs for Read access
  write:
    - '' # Insert ARNs...
  list:
    - '' # Insert ARNs...
  tag:
    - '' # Insert ARNs...
  permissions-management:
    - '' # Insert ARNs...
```

Then just fill it out:

```
policy_with_crud_levels:
- name: myRole
  description: 'Justification for privileges'
  arn: 'arn:aws:iam::123456789012:role/myRole'
  read:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  write:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  list:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  tag:
    - 'arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret'
```

(continues on next page)

(continued from previous page)

permissions-management:

```
- 'arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret'
```

Then run this command:

```
policy_sentry write-policy --crud --input-file crud.yml
```

It will generate these results:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "SsmReadParameter",
      "Effect": "Allow",
      "Action": [
        "ssm:getparameter",
        "ssm:getparameterhistory",
        "ssm:getparameters",
        "ssm:getparametersbypath",
        "ssm:listtagsforresource"
      ],
      "Resource": [
        "arn:aws:ssm:us-east-1:123456789012:parameter/myparameter"
      ]
    },
    {
      "Sid": "SsmWriteParameter",
      "Effect": "Allow",
      "Action": [
        "ssm:deleteparameter",
        "ssm:deleteparameters",
        "ssm:putparameter",
        "ssm:labelparameterversion"
      ],
      "Resource": [
        "arn:aws:ssm:us-east-1:123456789012:parameter/myparameter"
      ]
    },
    {
      "Sid": "SecretsmanagerPermissionsmanagementSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:deleteresourcepolicy",
        "secretsmanager:putresourcepolicy"
      ],
      "Resource": [
        "arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret"
      ]
    },
    {
      "Sid": "SecretsmanagerTaggingSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:tagresource",
        "secretsmanager:untagresource"
      ],
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```
        "Resource": [
            "arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret"
        ]
    }
]
}
```

Notice how the policy above recognizes the ARNs that the user supplies, along with the requested access level. For instance, the SID “SecretsmanagerTaggingSecret” contains Tagging actions that are assigned to the secret resource type only.

This rapidly speeds up the time to develop IAM policies, and ensures that all policies created limit access to exactly what your role needs access to. This way, developers only have to determine the resources that they need to access, and we abstract the complexity of IAM policies away from their development processes.

1.2 Installation

- `policy_sentry` is available via `pip`. To install, run:

```
pip install --user policy_sentry
```

1.2.1 Usage

- `initialize`: Create a SQLite database that contains all of the services available through the [Actions, Resources, and Condition Keys](#) documentation. See the [documentation](#).
- `create-template`: Creates the YAML file templates for use in the `write-policy` command types.
- `write-policy`: Leverage a YAML file to write policies for you
 - Option 1: Specify CRUD levels (Read, Write, List, Tagging, or Permissions management) and the ARN of the resource. It will write this for you. See the [documentation on CRUD mode](#)
 - Option 2: Specify a list of actions. It will write the IAM Policy for you, but you will have to fill in the ARNs. See the [documentation on Action Mode](#).
- `write-policy-dir`: This can be helpful in the Terraform use case. For more information, see the [wiki](#).
- `download-policies`: Download IAM policies from your AWS account for analysis.
- `analyze-iam-policy`: Analyze an IAM policy read from a JSON file, expands the wildcards (like `s3:List*` if necessary).
 - Option 1: Audits them to see if certain IAM actions are permitted, based on actions in a separate text file. See the [documentation on Initialization](#).
 - Option 2: Audits them to see if any of the actions in the policy meet a certain access level, such as “Permissions management.”

1.3 Author Information

Author:

- [Kinnaird McQuade](#)

- Twitter
- Keybase
- LinkedIn

Contributors:

- Matt Jones
 - Twitter
 - Keybase
 - LinkedIn

Comparison to other tools

2.1 Policy Revocation Tools

2.1.1 Repokid

RepoKid is a popular tool that was developed by Netflix, and is one of the more mature and battle-tested AWS IAM open source projects. It leverages AWS Access Advisor, which informs you how many AWS services your IAM Principal has access to, and how many of those services it has used in the last X amount of days or months. If you haven't used a service within the last 30 days, it "repos" your policy, and strips it of the privileges it doesn't use. It has some advanced features to allow for whitelisting roles and overall is a great tool.

One shortcoming is that AWS IAM Access Advisor only provides details at the service level (ex: S3-wide, or EC2-wide) and not down to the IAM Action level, so the revised policy is not very granular. However, RepoKid plays a unique role in the IAM ecosystem right now in that there are not any open source tools that provide similar functionality. **For that reason, it is best to view RepoKid and Policy Sentry as complimentary.**

Travis McPeak summarized the potential dynamic between Policy Sentry and RepoKid very well on [Clint Glibert's blog](#):

Policy Sentry aims to make it easy to create initial least privilege policies and then Repokid takes away unused permissions.

Creating policies is difficult, so Policy Sentry creates policies based on top level goals and target resources, and then on the backend substitutes the applicable action list to generate the policy. This is very helpful for anybody creating the first version of a policy.

To help with simplicity these permissions will be assigned somewhat coarsely. So Repokid can use data to remove the specific actions that were granted and aren't required. Also Repokid will repo down unused permissions once an application stops being used or scope changes.

2.2 AWS Tools

2.2.1 AWS Console - Visual Policy Editor


- [AWS IAM Visual Policy Editor in the AWS Console](#)

This policy generator is great in general and you should use it if you're getting used to writing IAM policies.

It's very similar to `policy_sentry` - you are able to bulk select according to access levels.

However, there are a number of downsides:

- **Missing access level type:** It does not specifically flag "Permissions management" access level
- **No override capabilities for inaccurate Access Levels:** Note how the `ssm:PutParameter` action is listed as "Tagging". This is inaccurate; it should be "Write". `Policy_sentry` allows you to override inaccurate access levels, whereas the Visual Policy Editor has had inaccurate Access levels for the last several years without any fixes.



- **Not automated:** Policy Sentry is, by design, meant for automated policy generation, whereas the Visual Policy Editor is meant to be manual.
- **Console Access:** It also requires access to the AWS Console.
- **Extensibility:** It's open source and Pull Requests are welcome! With `policy_sentry`, we get more control.

On the positive side, it does walk you through creating policies with IAM Condition keys. However, we believe that `policy_sentry`'s approach, where we **always** have policies restricted to the least amount of resources - provides a greater benefit to the end user. Furthermore, we plan on supporting condition keys at some point in the future.

2.2.2 AWS Policy Generator (static website)

- [AWS Policy Generator - static website](#)

AWS Policy Generator is a great tool; it supports IAM policies, as well as multiple types of resource-based policies (SQS Queue policy, S3 bucket policy, SNS Topic Policy, and VPC Endpoint Policy).

Loose ARN formatting: The regex expressions that it uses per-service does not require that actual valid resource ARNs are met - just that they meet the Regex requirement, which is uniform per-service. It just isn't as accurate or up to date as the actual IAM policy generation through the AWS Console

Missing actions: To determine the list of actions, it relies on a file titled `policies.js`, which contains a list of IAM Actions. However, this file is not as well maintained as the [Actions](#), [Resources](#), and [Condition Keys](#) tables. For example, it does not have these actions:

```
a4b:describe*
appstream:get*
cloudformation:preview*
codestar:verify*
ds:check*
health:get*
health:list*
kinesisanalytics:get*
lightsail:list*
mobilehub:validate*
resource-groups:describe*
```

2.3 Log-based Policy Generators

2.3.1 CloudTracker

- [CloudTracker](#)

Policy Sentry is somewhat similar to CloudTracker. CloudTracker queries CloudTrail logs using Amazon Athena and attempts to “guess” the matching between CloudTrail actions and IAM actions, then generates a policy. Given that there is not a 1-to-1 mapping between the names of Actions listed in CloudTrail log entries and the names AWS IAM Actions, the results are not always accurate. It is a good place to start, but the generated policies all contain Resources: “*”, so it is up to the user to restrict those IAM actions to only the necessary resources.

2.3.2 Trailscraper

- [Trailscraper](#)

Trailscraper does automated policy generation from CloudTrail logs, but there are some major limitations:

1. The generated policies have Resources set to ‘*’, not to a specific resource ARN
2. It downloads all of the CloudTrail logs. This takes a while.
 - Cloudtracker (<https://github.com/duo-labs/cloudtracker>) uses Amazon Athena, which is more efficient. In the future, I’d like to see a combined approach between all three of these tools to generate IAM policies based on Cloudtrail logs.
3. It is accurate to the point where there is a 1-to-1 mapping with the IAM actions vs CloudTrail logs. As I mentioned in other comments, since not every IAM Action is logged in CloudTrail and not every CloudTrail action matches IAM Actions, the results are not always accurate.

2.4 Other Infrastructure as Code Tools

2.4.1 aws-iam-generator

- [aws-iam-generator](#)

aws-iam-generator still requires you to write the actual policy templates from scratch, and then they allow you to re-use those policy templates.

Consider the JSON under [this area](#) of their README.

It’s essentially a method for managing their policies as code - but it doesn’t make those policies restricted to certain resources, unless you configure it that way. Using `policy_sentry --write-policy --crud`, you have to

supply a file with resource ARNs, and it will write the policy for you, rather than supplying a policy file, and hoping the ARNs fit that use case.

2.4.2 Terraform

The rationale described above also generally applies to Terraform, in that it still requires you to write the actual policy templates from scratch, and then you can re-use those policy templates. However, you still need to make those policies secure by default.

CHAPTER 3

Installation

- `policy_sentry` is available via `pip` (Python 3 only). To install, run:

```
pip install --user policy_sentry
```


CHAPTER 4

Initialization

initialize: This will create a SQLite database that contains all of the services available through the [Actions](#), [Resources](#), and [Condition Keys](#) documentation.

The database is stored in `$HOME/.policy_sentry/aws.sqlite3`.

The database is generated based on the HTML files stored in the `policy_sentry/shared/data/docs/` directory.

4.1 Options

- `--access-level-overrides-file` (Optional): Path to your own custom access level overrides file, used to override the Access Levels per action provided by AWS docs. The default one is [here](#).
- `--fetch` (Optional): Specify this flag to fetch the HTML Docs directly from the AWS website. This will be helpful if the docs in the Git repository are behind the live docs and you need to use the latest version of the docs right now.
- `--fetch` (Optional) Specify this flag to fetch the HTML Docs directly from the AWS website. This will be helpful if the docs in the Git repository are behind the live docs and you need to use the latest version of the docs right now.

4.2 Usage

```
# Initialize the database, using the existing Access Level Overrides file
policy_sentry initialize

# Fetch the most recent version of the AWS documentation so you can experiment with
↳ new services.
# This can be helpful in case the AWS HTML files in the Python package are outdated,
↳ even if it is a week old
```

(continues on next page)

(continued from previous page)

```
policy_sentry initialize --fetch  
  
# Initialize the database with a custom Access Level Overrides file  
  
policy_sentry initialize --access-level-overrides-file ~/.policy_sentry/access-level-  
↳overrides.yml  
policy_sentry initialize --access-level-overrides-file ~/.policy_sentry/overrides-  
↳resource-policies.yml
```

4.3 Skipping Initialization

When using Policy Sentry manually, you have to build a local database file with the initialize function.

However, if you are developing your own Python code and you want to import Policy Sentry as a third party package, you can skip the initialization and leverage the local database file that is bundled with the Python package itself.

This is especially useful for developers who wish to leverage Policy Sentry's capabilities that require the use of the IAM database (such as querying the IAM database table). This way, you don't have to initialize the database and can just query it immediately.

5.1 CRUD Mode: ARNs and Access Levels

This is the flagship feature of this tool. You can just specify the CRUD levels (Read, Write, List, Tagging, or Permissions management) for each action in a YAML File. The policy will be generated for you. You might need to fiddle with the results for your use in Terraform, but it significantly reduces the level of effort to build least privilege into your policies.

5.1.1 Command options

- `--crud`: Specify this option to use the CRUD functionality. File must be formatted as expected. Defaults to false.
- `--input-file`: YAML file containing the CRUD levels + Resource ARNs. Required.
- `--minimize`: Whether or not to minimize the resulting statement with *safe* usage of wildcards to reduce policy length. Set this to the character length you want. This can be extended for readability. I suggest setting it to 0.

Example:

```
policy_sentry write-policy --crud --input-file examples/crud.yml
```

5.1.2 Instructions

- To generate a policy according to resources and access levels, start by creating a template with this command so you can just fill out the ARNs:

```
policy_sentry create-template --name myRole --output-file crud.yml --template-type_  
↪crud
```

- It will generate a file like this:

```
policy_with_crud_levels:
- name: myRole
  description: ''
  arn: ''
  # Insert ARNs below
  read:
    - ''
  write:
    - ''
  list:
    - ''
  tag:
    - ''
  permissions-management:
    - ''
  # Provide a list of IAM actions that cannot be restricted to ARNs
  wildcard:
    - ''
```

- Then just fill it out:

```
policy_with_crud_levels:
- name: myRole
  description: 'Justification for privileges'
  arn: 'arn:aws:iam::123456789102:role/myRole'
  read:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  write:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  list:
    - 'arn:aws:ssm:us-east-1:123456789012:parameter/myparameter'
  tag:
    - 'arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret'
  permissions-management:
    - 'arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret'
```

- Run the command:

```
policy_sentry write-policy --crud --input-file examples/crud.yml
```

- It will generate an IAM Policy containing an IAM policy with the actions restricted to the ARNs specified above.
- The resulting policy (without the `--minimize` command) will look like this:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "SsmReadParameter",
      "Effect": "Allow",
      "Action": [
        "ssm:getparameter",
        "ssm:getparameterhistory",
        "ssm:getparameters",
        "ssm:getparametersbypath",
        "ssm:listtagsforresource"
      ],
      "Resource": [
```

(continues on next page)

(continued from previous page)

```

        "arn:aws:ssm:us-east-1:123456789012:parameter/myparameter"
    ],
    },
    {
        "Sid": "SsmWriteParameter",
        "Effect": "Allow",
        "Action": [
            "ssm:deleteparameter",
            "ssm:deleteparameters",
            "ssm:putparameter",
            "ssm:labelparameterversion"
        ],
        "Resource": [
            "arn:aws:ssm:us-east-1:123456789012:parameter/myparameter"
        ]
    },
    {
        "Sid": "SecretsmanagerPermissionsmanagementSecret",
        "Effect": "Allow",
        "Action": [
            "secretsmanager:deleteresourcepolicy",
            "secretsmanager:putresourcepolicy"
        ],
        "Resource": [
            "arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret"
        ]
    },
    {
        "Sid": "SecretsmanagerTaggingSecret",
        "Effect": "Allow",
        "Action": [
            "secretsmanager:tagresource",
            "secretsmanager:untagresource"
        ],
        "Resource": [
            "arn:aws:secretsmanager:us-east-1:123456789012:secret:mysecret"
        ]
    }
]
}

```

5.2 Actions Mode: Lists of IAM Actions

Supply a list of actions in a YAML file and generate the policy accordingly.

5.2.1 Command options

- `--input-file`: YAML file containing the list of actions
- `--minimize`: Whether or not to minimize the resulting statement with *safe* usage of wildcards to reduce policy length. Set this to the character length you want - for example, 4

Example:

```
policy_sentry write-policy --input-file examples/actions.yml
```

5.2.2 Instructions

- If you already know the IAM actions, you can just run this command to create a template to fill out:

```
policy_sentry create-template --name myRole --output-file tmp.yml --template-type_↵
↵actions
```

- It will generate a file with contents like this:

```
policy_with_actions:
- name: myRole
  description: '' # Insert value here
  arn: '' # Insert value here
  actions:
- '' # Fill in your IAM actions here
```

- Create a yaml file with the following contents:

```
policy_with_actions:
- name: 'RoleNameWithActions'
  description: 'Justification for privileges' # for auditability
  arn: 'arn:aws:iam::123456789102:role/myRole' # for auditability
  actions:
  - kms:CreateGrant
  - kms:CreateCustomKeyStore
  - ec2:AuthorizeSecurityGroupEgress
  - ec2:AuthorizeSecurityGroupIngress
```

- Then run this command:

```
policy_sentry write-policy --input-file examples/actions.yml
```

- The output will look like this:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "KmsPermissionsmanagementKey",
      "Effect": "Allow",
      "Action": [
        "kms:creategrant"
      ],
      "Resource": [
        "arn:aws:kms:${Region}:${Account}:key/${KeyId}"
      ]
    },
    {
      "Sid": "Ec2WriteSecuritygroup",
      "Effect": "Allow",
      "Action": [
        "ec2:authorizesecuritygroupegress",
        "ec2:authorizesecuritygroupingress"
      ]
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "Resource": [
        "arn:aws:ec2:${Region}:${Account}:security-group/${SecurityGroupId}"
    ]
},
{
    "Sid": "MultMultNone",
    "Effect": "Allow",
    "Action": [
        "kms:createcustomkeystore",
        "cloudhsm:describeclusters"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

5.3 Folder Mode: Write Multiple Policies from CRUD mode files

This command provides the same function as *write-policy*'s CRUD mode, but it can execute all the CRUD mode files in a folder. This is particularly useful in the Terraform use case, where the Terraform module can export a number of Policy Sentry template files into a folder, which can then be consumed using this command.

See the Terraform demo for more details.

```
Usage: policy_sentry write-policy-dir [OPTIONS]
```

Options:

```

--input-dir TEXT      Relative path to Input directory that contains policy_sentry .
↳ yml files (CRUD mode only)  [required]
--output-dir TEXT     Relative path to directory to store AWS JSON policies [required]
--crud               Use the CRUD functionality. Defaults to false
--minimize INTEGER    Minimize the resulting statement with *safe* usage of wildcards,
↳ to reduce policy length. Set this to the character length you want - for example, 4
--help               Show this message and exit.

```

Downloading Policies

```
Usage: policy_sentry download-policies [OPTIONS]
```

```
Download remote IAM policies to a directory for use in the analyze-iam-  
policies command.
```

```
Options:
```

```
--recursive          Use this flag to download *all* IAM policies from  
                      accounts listed in your AWS credentials file.  
--profile TEXT        To authenticate to AWS and analyze *all* existing IAM  
                      policies.  
--aws-managed         Use flag if you want to download AWS Managed policies  
                      too.  
--include-unattached Download both attached and unattached policies.  
--help               Show this message and exit.
```

- Make sure you are authenticated to AWS.

6.1 Customer-managed policies - one account

- Run this command:

```
policy_sentry download-policies --profile dev
```

- It will download the policies to `$HOME/.policy_sentry/policy-analysis/account-number/customer-managed`.
- You can then run analysis on the entire directory:

```
policy_sentry analyze
```

Then it will generate a report based on risky IAM actions for a variety of categories, like Network Exposure, Resource Exposure, Credentials Exposure, or Privilege Escalation.

6.2 AWS Managed policies

- Run this command:

```
policy_sentry download-policies --profile dev --aws-managed
```

- It will download the policies to `$HOME/.policy_sentry/policy-analysis/account-number/aws-managed`.
- You can then run analysis on the entire directory:

```
analyze-iam-policy --policy $HOME/.policy_sentry/policy-analysis/0123456789012/  
↪customer-managed --from-access-level permissions-management
```

Then it will print out the AWS Managed IAM policies that contain actions with “Permissions management” access levels.

Analyzing Policies

`analyze-iam-policy`: Reads policies downloaded locally, and expands the wildcards (like `s3:List*` if necessary, and audits them to see if certain IAM actions are permitted.

7.1 Motivation

Case 1:

Let's say that you have hundreds or thousands of engineers at your organization who have permissions to create or edit IAM policies. Perhaps this is just in Dev, or it's just allowed via Infrastructure as Code, or maybe there is rampant shadow IT.

You know that there is a huge issue with overpermissioning at your organization - not only is there widespread use of `Resources: "*" in IAM policies`, lots of engineers have the ability to create network resources, update security groups, create resource-based policies, escalate privileges via methods popularized through [Rhino Security Labs research](#), and more. To convince your boss that a serious IAM uplift project should happen, you want to show your boss **the inherent risk levels in IAM policies per account, regardless of whether or not it is currently a vulnerability**.

If you're in this position - welcome to the right place.

Case 2:

Let's say you are a developer that handles creation of IAM policies.

- An internal customer asks you to create an IAM policy.
- You haven't been tasked with auditing IAM policies yourself, as that's not your area of expertise, and until this point there is no automation to do it for you.
- However, you want to make sure that the customers aren't asking for permissions that they don't need, since we need to have *some* guardrails in place to prevent unnecessary exposure of attack surfaces.
- This is made more difficult by the fact that sometimes, the customer will give you IAM policies that include `*` in the actions. Not only do you want to restrict actions to the specific ARNs, but you want to know what actions they actually need!

You can solve this with `policy_sentry` too, by auditing for IAM actions in a given policy. Tell them to supply the policy to you in JSON format, stash it in `~/.policy_sentry/analysis/account_id/customer-managed/`, and feed it into the `analyze_iam_policy` command, as shown below.

7.2 Options

```
Usage: policy_sentry analyze [OPTIONS] COMMAND [ARGS]...
```

Analyze locally stored IAM policies and generate a report.

Options:

`--help` Show this message and exit.

Commands:

<code>downloaded-policies</code>	Analyze <i>*all*</i> locally downloaded IAM policy files and generate a report.
<code>policy-file</code>	Analyze a <i>*single*</i> policy file and generate a report

downloaded-policies subcommand:

```
Usage: policy_sentry analyze downloaded-policies [OPTIONS]
```

Analyze all locally downloaded IAM policy files and generate a report.

Options:

<code>--report-config PATH</code>	Custom report configuration file. Contains policy name exclusions and custom risk score weighting. Defaults to <code>~/.policy_sentry/report-config.yml</code>
<code>--report-name TEXT</code>	Name of the report. Defaults to "overall".
<code>--include-markdown-report</code>	Use this flag to enable a Markdown report, which can be used with pandoc to generate an HTML report. Due to potentially very large report sizes, this is set to False by default.
<code>--help</code>	Show this message and exit.

policy-file subcommand:

```
Usage: policy_sentry analyze policy-file [OPTIONS]
```

Analyze a **single** policy file and generate a report

Options:

<code>--policy PATH</code>	The policy file to analyze. [required]
<code>--report-config PATH</code>	Custom report configuration file. Contains policy name exclusions and custom risk score weighting. Defaults to <code>~/.policy_sentry/report-config.yml</code>
<code>--report-path PATH</code>	<i>*Path*</i> to the directory of the final report. Defaults to current directory.
<code>--account-id TEXT</code>	Account ID for the policy. If you want the report to include the account ID, provide it here. Defaults to a placeholder value.
<code>--include-markdown-report</code>	Use this flag to enable a Markdown report, which can be used with pandoc to generate an HTML report. Due to potentially very large report sizes, this is set to False by default.
<code>--help</code>	Show this message and exit.

7.3 Instructions

- Don't forget to build the database first:

```
policy_sentry initialize
```

7.3.1 Risk Categories

1. **Privilege Escalation:** This is based off of [Rhino Security Labs research](#)
2. **Resource Exposure:** This contains all IAM Actions at the “Permissions Management” resource level. Essentially - if your policy can (1) write IAM Trust Policies, (2) write to the RAM service, or (3) write Resource-based Policies, then the action has the potential to result in resource exposure if an IAM principal with that policy was compromised.
3. **Network Exposure:** This highlights IAM actions that indicate an IAM principal possessing these actions could create resources that could be exposed to the public at the network level. For example, public RDS clusters, public EC2 instances. While possession of these privileges does not constitute a security vulnerability, it is important to know exactly who has these permissions.
4. **Credentials Exposure:** This includes IAM actions that grant some kind of credential, where if exposed, it could grant access to sensitive information. For example, `ecr:GetAuthorizationToken` creates a token that is valid for 12 hours, which you can use to authenticate to Elastic Container Registries and download Docker images that are private to the account.

7.3.2 Audit all downloaded policies and generate a report

- Command:

```
# 1. Use a tool like Gossamer (https://github.com/GESkunkworks/gossamer) to update_
↳ your AWS credentials profile all at once
# 2. Recursively download all IAM policies from accounts in your credentials file
# Note: alternatively, you can just place them there yourself.
policy_sentry download --recursive

# Audit all JSON policies under the path ~/.policy_sentry/analysis/account_id/
↳ customer-managed
policy_sentry analyze --downloaded-policies

# Use a custom report configuration. This is typically used for excluding role names.
↳ Defaults to ~/.policy_sentry/report-config.yml
policy_sentry analyze --downloaded-policies --report-config custom-config.yml
```

- Output:

```
Analyzing...
/Users/kmcquade/.policy_sentry/analysis/0123456789012/
/Users/kmcquade/.policy_sentry/analysis/9876543210123/
...

Reports saved to:
-/Users/kmcquade/.policy_sentry/analysis/overall.json
-/Users/kmcquade/.policy_sentry/analysis/overall.csv

The JSON Report contains the raw data. The CSV report shows a report summary.
```

- The raw JSON data will look like this:

```
{
  "some-risky-policy": {
    "account_id": "0123456789012",
    "resource_exposure": [
      "iam:createaccesskey",
      "iam:deleteaccesskey"
    ],
    "privilege_escalation": [
      "iam:createaccesskey"
    ]
  },
  "another-risky-policy": {
    "account_id": "9876543210123",
    "resource_exposure": [
      "iam:updateassumerolepolicy",
      "iam:updaterole"
    ],
    "privilege_escalation": [
      "iam:updateassumerolepolicy"
    ],
    "credentials_exposure": [
      "ecr:getauthorizationtoken"
    ],
    "network_exposure": [
      "ec2:authorizesecuritygroupingress",
      "ec2:authorizesecuritygroupegress"
    ]
  }
}
```

7.3.3 Audit a single IAM policy and generate a report

- Command:

```
# Analyze a single IAM policy
policy_sentry analyze policy-file --policy examples/explicit-actions.json
```

- This will create a CSV file that looks like this:

Account ID	Policy Name	Resource Expo- sure	Privilege Escala- tion	Network Expo- sure	Credentials Expo- sure
000000000000	explicit-actions	9	0	0	1

- ... and a JSON data file that looks like this:

```
{
  "explicit-actions": {
    "resource_exposure": [
      "ecr:setrepositorypolicy",
      "s3:deletebucketpolicy",
      "s3:objectowneroverridetobucketowner",
      "s3:putaccountpublicaccessblock",
      "s3:putbucketacl",

```

(continues on next page)

(continued from previous page)

```
        "s3:putbucketpolicy",
        "s3:putbucketpublicaccessblock",
        "s3:putobjectacl",
        "s3:putobjectversionacl"
    ],
    "account_id": "000000000000",
    "credentials_exposure": [
        "ecr:getauthorizationtoken"
    ]
}
```

7.3.4 Custom Config file

- Quite often, organizations may have customer-managed policies that are in every account, or are very permissive by design. Rather than having a very large report every time you run this tool, you can specify a custom config file with this command. Just make sure you format it correctly, as shown below.

```
report-config:
  excluded-role-patterns:
    - "Administrator*
```

Note: This probably will eventually support: - Action-specific exclusions per-account and per-role - Turning risk categories on and off

Querying the Policy Database

Policy Sentry relies on a SQLite database, generated at *initialize* time, which contains all of the services available through the [Actions, Resources, and Condition Keys](#) documentation. The HTML files from that AWS documentation is scraped and stored in the SQLite database, which is then stored in `$HOME/.policy_sentry/aws.sqlite3`.

Policy Sentry supports querying that database through the CLI. This can help with writing policies and generally knowing what values to supply in your policies.

8.1 Commands

- Query the **Action** table:

```
# Get a list of all IAM Actions available to the RAM service
policy_sentry query action-table --service ram

# Get details about the `ram:TagResource` IAM Action
policy_sentry query action-table --service ram --name tagresource

# Get a list of all IAM actions under the RAM service that have the Permissions_
↳management access level.
policy_sentry query action-table --service ram --access-level permissions-management

# Get a list of all IAM actions under the SES service that support the_
↳`ses:FeedbackAddress` condition key.
policy_sentry query action-table --service ses --condition ses:FeedbackAddress
```

- Query the **ARN** table:

```
# Get a list of all RAW ARN formats available through the SSM service.
policy_sentry query arn-table --service ssm

# Get the raw ARN format for the `cloud9` ARN with the short name `environment`
policy_sentry query arn-table --service cloud9 --name environment
```

(continues on next page)

(continued from previous page)

```
# Get key/value pairs of all RAW ARN formats plus their short names
policy_sentry query arn-table --service cloud9 --list-arn-types
```

- Query the **Condition Keys** table:

```
# Get a list of all condition keys available to the Cloud9 service
policy_sentry query condition-table --service cloud9

# Get details on the condition key titled `cloud9:Permissions`
policy_sentry query condition-table --service cloud9 --name cloud9:Permissions
```

8.2 Options

- action-table

```
Usage: policy_sentry query action-table [OPTIONS]

Options:
  --service TEXT      Filter according to AWS service. [required]
  --name TEXT         The name of IAM Action. For example, if the
                      action is "iam:ListUsers", supply
                      "ListUsers" here.
  --access-level [read|write|list|tagging|permissions-management]
                      If action table is chosen, you can use this
                      to filter according to CRUD levels.
                      Acceptable values are read, write, list,
                      tagging, permissions-management
  --condition TEXT    If action table is chosen, you can supply a
                      condition key to show a list of all IAM
                      actions that support the condition key.
  --wildcard-only      If action table is chosen, show the IAM
                      actions that only support wildcard resources
                      - i.e., cannot support ARNs in the resource
                      block.
  --help              Show this message and exit.
```

- arn-table

```
Usage: policy_sentry query arn-table [OPTIONS]

Options:
  --service TEXT      Filter according to AWS service. [required]
  --name TEXT         The short name of the resource ARN type. For example,
                      `bucket` under service `s3`.
  --list-arn-types    If ARN table is chosen, show the short names of ARN Types.
                      If empty, this will show RAW ARNs only.
  --help              Show this message and exit.
```

- condition-table

```
Usage: policy_sentry query condition-table [OPTIONS]

Options:
```

(continues on next page)

(continued from previous page)

<code>--name TEXT</code>	Get details on a specific condition key. Leave this blank to get a list of all condition keys available to the service.
<code>--service TEXT</code>	Filter according to AWS service. [required]
<code>--help</code>	Show this message and exit.

CHAPTER 9

Docker

If you prefer using Docker instead of installing the script with Python, we support that as well.

Use this to build the docker image:

```
docker build -t kmcquade/policy_sentry .
```

Use this to run some basic commands:

```
# Basic commands with no arguments
docker run -i --rm kmcquade/policy_sentry:latest "--help"
docker run -i --rm kmcquade/policy_sentry:latest "query"

# Query the database
docker run -i --rm kmcquade/policy_sentry:latest "query action-table --service all --
↳access-level permissions-management"
```

The *write-policy* command also supports passing in the YAML config via STDIN. Try it out here:

```
# Write policies by passing in the config via STDIN
cat examples/yml/crud.yml | docker run -i --rm kmcquade/policy_sentry:latest "write-
↳policy --crud"
cat examples/yml/actions.yml | docker run -i --rm kmcquade/policy_sentry:latest
↳"write-policy"
```


10.1 Commands

- `initialize`: Create a SQLite database that contains all of the services available through the [Actions, Resources, and Condition Keys](#) documentation. See the [documentation](#).
- `download-policies`: Download IAM policies from an AWS account locally for analysis against the database.
- `create-template`: Creates the YAML file templates for use in the `write-policy` command types.
- `write-policy`: Leverage a YAML file to write policies for you
 - Option 1: **CRUD Mode**. Specify CRUD levels (Read, Write, List, Tagging, or Permissions management) and the ARN of the resource. It will write this for you. See the documentation for more details.
 - Option 2: **Actions Mode**. Specify a list of actions. It will write the IAM Policy for you, but you will have to fill in the ARNs. See the documentation for more details.
- `write-policy-dir`: This can be helpful in the Terraform use case. For more information, see the wiki.
- `analyze`: Analyze IAM policies downloaded locally, expands the wildcards (like `s3:List*`) if necessary, and generates a report based on policies that are flagged for these risk categories:
 1. **Privilege Escalation**: This is based off of [Rhino Security Labs](#) research
 2. **Resource Exposure**: This contains all IAM Actions at the “Permissions Management” resource level. Essentially - if your policy can (1) write IAM Trust Policies, (2) write to the RAM service, or (3) write Resource-based Policies, then the action has the potential to result in resource exposure if an IAM principal with that policy was compromised.
 3. **Network Exposure**: This highlights IAM actions that indicate an IAM principal possessing these actions could create resources that could be exposed to the public at the network level. For example, public RDS clusters, public EC2 instances. While possession of these privileges does not constitute a security vulnerability, it is important to know exactly who has these permissions.
 4. **Credentials Exposure**: This includes IAM actions that grant some kind of credential, where if exposed, it could grant access to sensitive information. For example, `ecr:GetAuthorizationToken` creates

a token that is valid for 12 hours, which you can use to authenticate to Elastic Container Registries and download Docker images that are private to the account.

- query: Query the IAM database tables. This can help when filling out the Policy Sentry templates, or just querying the database for quick knowledge.
 - Option 1: Query the Actions Table (`--table action`)
 - Option 2: Query the ARNs Table (`--table arn`)
 - Option 3: Query the Conditions Table (`--table condition`)

10.2 Initialization

```
# Initialize the policy_sentry config folder and create the IAM database tables.
policy_sentry initialize

# Fetch the most recent version of the AWS documentation so you can experiment with
↳ new services.
policy_sentry initialize --fetch

# Override the Access Levels by specifying your own Access Levels (example:,
↳ correcting Permissions management levels)
policy_sentry initialize --access-level-overrides-file ~/.policy_sentry/access-level-
↳ overrides.yml
policy_sentry initialize --access-level-overrides-file ~/.policy_sentry/overrides-
↳ resource-policies.yml
```

10.3 Policy Writing Commands

```
# Initialize the policy_sentry config folder and create the IAM database tables.
policy_sentry initialize

# Create a template file for use in the write-policy command (crud mode)
policy_sentry create-template --name myRole --output-file tmp.yml --template-type crud

# Write policy based on resource-specific access levels
policy_sentry write-policy --crud --input-file examples/yml/crud.yml

# Write policy_sentry YAML files based on resource-specific access levels on a
↳ directory basis
policy_sentry write-policy-dir --crud --input-dir examples/input-dir --output-dir
↳ examples/output-dir

# Create a template file for use in the write-policy command (actions mode)
policy_sentry create-template --name myRole --output-file tmp.yml --template-type
↳ actions

# Write policy based on a list of actions
policy_sentry write-policy --input-file examples/yml/actions.yml
```

10.4 IAM Database Query Commands

- Query the **Action** table:

```
# Get a list of all IAM actions across ALL services that have "Permissions management
↳" access
policy_sentry query action-table --service all --access-level permissions-management

# Get a list of all IAM Actions available to the RAM service
policy_sentry query action-table --service ram

# Get details about the `ram:TagResource` IAM Action
policy_sentry query action-table --service ram --name tagresource

# Get a list of all IAM actions under the RAM service that have the Permissions_
↳management access level.
policy_sentry query action-table --service ram --access-level permissions-management

# Get a list of all IAM actions under the SES service that support the_
↳`ses:FeedbackAddress` condition key.
policy_sentry query action-table --service ses --condition ses:FeedbackAddress
```

- Query the **ARN** table:

```
# Get a list of all RAW ARN formats available through the SSM service.
policy_sentry query arn-table --service ssm

# Get the raw ARN format for the `cloud9` ARN with the short name `environment`
policy_sentry query arn-table --service cloud9 --name environment

# Get key/value pairs of all RAW ARN formats plus their short names
policy_sentry query arn-table --service cloud9 --list-arn-types
```

- Query the **Condition Keys** table:

```
# Get a list of all condition keys available to the Cloud9 service
policy_sentry query condition-table --service cloud9
# Get details on the condition key titled `cloud9:Permissions`
policy_sentry query condition-table --service cloud9 --name cloud9:Permissions
```

10.5 Policy Download and Analysis Commands

```
# Initialize the policy_sentry config folder and create the IAM database tables.
policy_sentry initialize

# Download customer managed IAM policies from a live account under 'default' profile.
↳By default, it looks for policies that are 1. in use and 2. customer managed
policy_sentry download-policies # this will download to ~/.policy_sentry/accountid/
↳customer-managed/.json

# Download customer-managed IAM policies, including those that are not attached
policy_sentry download-policies --include-unattached # this will download to ~/.
↳policy_sentry/accountid/customer-managed/*.json
```

(continues on next page)

(continued from previous page)

```
# Analyze a single IAM policy FILE
policy_sentry analyze policy-file --policy examples/explicit-actions.json

# 1. Use a tool like Gossamer (https://github.com/GESkunkworks/gossamer) to update_
↳ your AWS credentials profile all at once
# 2. Recursively download all IAM policies from accounts in your credentials file
policy_sentry download-policies --recursive

# Audit all IAM policies downloaded locally and generate CSV and JSON reports.
policy_sentry analyze downloaded-policies

# Audit all IAM policies and also include a Markdown formatted report, then convert_
↳ it to HTML
policy_sentry analyze --include-markdown-report
pandoc -f markdown ~/.policy_sentry/analysis/overall.md -t html > overall.html

# Use a custom report configuration. This is typically used for excluding role names._
↳ Defaults to ~/.policy_sentry/report-config.yml
policy_sentry analyze --report-config custom-config.yml
```

Please download the demo code [here](#) to follow along.

11.1 Command options

```
Usage: policy_sentry write-policy-dir [OPTIONS]

write_policy, but this time with an input directory of YML/YAML files, and
an output directory for all the JSON files

Options:
  --input-dir TEXT      Relative path to Input directory that contains policy_sentry .
                        ↪ yml files (CRUD mode only) [required]
  --output-dir TEXT     Relative path to directory to store AWS JSON policies [required]
  --crud               Use the CRUD functionality. Defaults to false
  --minimize INTEGER    Minimize the resulting statement with *safe* usage of wildcards,
                        ↪ to reduce policy length. Set this to the character length you want - for example, 4
  --help               Show this message and exit.
```

11.2 Prerequisites

This requires:

- Terraform v0.12.8
- AWS credentials; must be authenticated

11.3 Tutorial

- Install policy_sentry

```
pip3 install policy_sentry
```

- Initialize policy_sentry

```
policy_sentry initialize
```

- Execute the first Terraform module:

```
cd environments/standard-resources
tfjson install 0.12.8
terraform init
terraform plan
terraform apply -auto-approve
```

This will create a YAML file to be used by policy_sentry in the `environments/iam-resources/files/` directory titled `example-role-randomid.yml`.

- Write the policy using policy_sentry:

```
cd ../iam-resources
policy_sentry write-policy-dir --crud --input-dir files --output-dir files
```

This will create a JSON file to be consumed by Terraform's `aws_iam_policy` resource to create an IAM policy.

- Now create the policies with Terraform:

```
terraform init
terraform plan
terraform apply -auto-approve
```

- Don't forget to cleanup

```
terraform destroy -auto-approve
cd ../standard-resources
terraform destroy -auto-approve
```

12.1 1: Install policy_sentry

- Install policy_sentry

```
pip3 install policy_sentry
```

- Initialize policy_sentry

```
policy_sentry initialize
```

12.2 2: Generate the policy_sentry YAML File

Create a file with the following in some-directory:

```
module "policy_sentry_yaml" {
  source      = "git::https://github.com/salesforce/policy_sentry.git//examples/
↳ terraform/modules/generate-policy_sentry-yml"
  role_name   = ""
  role_description = ""
  role_arn    = ""

  list_access_level           = []
  permissions_management_access_level = []
  read_access_level          = []
  tagging_access_level       = []
  write_access_level         = []

  yaml_file_destination_path = "../other-directory/files"
}
```

Make sure you fill out the actual directory path properly. Note that `yml_file_destination_path` should point to the directory mentioned in Step 3.

12.3 3: Run policy_sentry and specify proper target directory

- Enter the directory you specified under `yml_file_destination_path` above.
- Run the following:

```
policy_sentry write-policy-dir --crud --input-dir files --output-dir files
```

12.4 4: Create the IAM Policies using JSON files from directory

Then from other-directory:

```
module "policies" {  
  source = "git::https://github.com/salesforce/policy_sentry.git//examples/terraform/  
↩modules/generate-iam-policies"  
  relative_path_to_json_policy_files = "files"  
}
```


Want to contribute back to Policy Sentry? This page describes the general development flow, our philosophy, the test suite, and issue tracking.

Impostor Syndrome Disclaimer

Before we get into the details: *We want your help. No, really.*

There may be a little voice inside your head that is telling you that you're not ready to be an open source contributor; that your skills aren't nearly good enough to contribute. What could you possibly offer a project like this one?

We assure you – the little voice in your head is wrong. If you can write code at all, you can contribute code to open source. Contributing to open source projects is a fantastic way to advance one's coding skills. Writing perfect code isn't the measure of a good developer (that would disqualify all of us!); it's trying to create something, making mistakes, and learning from those mistakes. That's how we all improve.

We've provided some clear Contribution Guidelines that you can read [here](#). The guidelines outline the process that you'll need to follow to get a patch merged. By making expectations and process explicit, we hope it will make it easier for you to contribute.

And you don't just have to write code. You can help out by writing documentation, tests, or even by giving feedback about this work. (And yes, that includes giving feedback about the contribution guidelines.)

([Adrienne Friend](#) came up with this disclaimer language.)

13.1 Contributing to Documentation

If you're looking to help document Policy Sentry, your first step is to get set up with Sphinx, our documentation tool. First you will want to make sure you have a few things on your local system:

- python-dev (if you're on OS X, you already have this)
- pip
- pipenv

Once you've got all that, the rest is simple:

```
# If you have a fork, you'll want to clone it instead
git clone git@github.com:salesforce/policy_sentry.git

# Set up the Pipenv
pipenv install --skip-lock
pipenv shell

# Enter the docs directory and compile
cd docs/
make html

# View the file titled docs/_build/html/index.html in your browser
```

13.1.1 Building Documentation

Inside the `docs` directory, you can run `make` to build the documentation. See `make help` for available options and the [Sphinx Documentation](#) for more information.

13.1.2 Docstrings

The comments under each Python Module are Docstrings. We use those to generate our documentation. See more information here: <https://sphinx-rtd-tutorial.readthedocs.io/en/latest/build-the-docs.html#generating-documentation-from-docstrings>.

Use the Google style for Docstrings, as shown here: <http://www.sphinx-doc.org/en/master/usage/extensions/napoleon.html#google-vs-numpy>

::

```
def func(arg1, arg2): """Summary line.
    Extended description of function.
    Args: arg1 (int): Description of arg1 arg2 (str): Description of arg2
    Returns: bool: Description of return value
    """ return True
```

13.2 IAM Database

Policy Sentry leverages HTML files from the [Actions, Resources, and Condition Keys](#) page AWS documentation to build the IAM database.

- These HTML files are included as part of the PyPi package
- The database itself is

This design choice was made for a few reasons:

1. **Don't break because of AWS:** The automation must **not** break if the AWS website is down, or if AWS drastically changes the documentation.
2. **Replicability:** Two `git clones` that build the SQLite database should always have the same results
3. **Easy to review:** The repository itself should contain easy-to-understand and easy-to-view documentation, which the user can replicate, to verify with the human eye that no malicious changes have been made.

- This means no JSON files with complicated structures, or Binary files (the latter of which does not permit `git diff`) in the repository.
- This helps to mitigate the concern that open source software could be modified to alter IAM permissions at other organizations.

13.2.1 How Policy Sentry uses the IAM database

policy_sentry follows this process for generating policies.

1. If **User-supplied actions** is chosen:
 - Look up the actions in our master Actions Table in the database, which contains the Action Tables for all AWS services
 - If the action in the database matches the actions requested by the user, determine the ARN Format required.
 - Proceed to step 3
2. If **User-supplied ARNs with Access levels** (i.e., the `--crud` flag) is chosen:
 - Match the user-supplied ARNs with ARN formats in our ARN Table database, which contains the ARN tables for all AWS Services
 - If it matches, get the access level requested by the user
 - Proceed to step 3
3. Compile those into groups, sorted by an SID namespace. The SID namespace follows the format of **Service**, **Access Level**, and **Resource ARN Type**, with no character delimiter (to meet AWS IAM Policy formatting expectations). For example, the namespace could be `SsmReadParameter`, `KmsReadKey`, or `Ec2TagInstance`.
4. Then, we associate the user-supplied ARNs matching that namespace with the SID.
5. If **User-supplied actions** is chosen:
 - Associate the IAM actions requested by the user to the service, access level, and ARN type matching the aforementioned SID namespace
6. If **User-supplied ARNs with Access levels** (i.e., the `--crud` flag) is chosen:
 - Associate all the IAM actions that correspond to the service, access level, and ARN type matching that SID namespace.
7. Print the policy

Updating the AWS HTML files

The command shown below downloads the Actions, Resources, and Condition Keys pages per-service to the `policy_sentry/shared/data/docs` folder.

- The HTML files will be stored in `policy_sentry/shared/data/docs/list_*.partial.html`
- It also add a file titled `policy_sentry/shared/data/links.yml` as well.
- It also builds a SQLite database file to include as part of the PyPi package.

This will update the

```
python3 ./utils/download_docs.py
```

This downloads the Actions, Resources, and Condition Keys pages per-service to the `policy_sentry/shared/data/docs` folder. It also add a file titled `policy_sentry/shared/data/links.yml` as well.

When a user runs `policy_sentry initialize`, these files are copied over to the config folder (`~/.policy_sentry/`).

This design choice was made for a few reasons:

1. **Don't break because of AWS:** The automation must **not** break if the AWS website is down, or if AWS drastically changes the documentation.
2. **Replicability:** Two *git clones* that build the SQLite database should always have the same results
3. **Easy to review: The repository itself should contain easy-to-understand and easy-to-view documentation, which the user can**
 - This means no JSON files with complicated structures, or Binary files (the latter of which does not permit `git diff`) in the repository.
 - This helps to mitigate the concern that open source software could be modified to alter IAM permissions at other organizations.

13.3 Testing

13.3.1 Pipenv

```
pipenv --python 3.7 # create the environment
pipenv shell       # start the environment
pipenv install     # install both development and production dependencies
```

13.3.2 Invoke

To run and develop Policy Sentry without having to install from PyPi, you can use Invoke.

```
# List available tasks
invoke -l

# that will show the following options:
Available tasks:

build.build-package      Build the policy_sentry package from the current
                           directory contents for use with PyPi
build.install-package    Install the policy_sentry package built from the
                           current directory contents (not PyPi)
build.uninstall-package  Uninstall the policy_sentry package
build.upload-prod        Upload the package to the PyPi production server
                           (requires credentials)
build.upload-test        Upload the package to the TestPyPi server
                           (requires credentials)
integration.analyze-policy Integration testing: Tests the `analyze`
                           functionality
integration.clean        Runs `rm -rf $HOME/.policy_sentry`
integration.initialize    Integration testing: Initialize the
                           policy_sentry database
integration.query        Integration testing: Tests the `query`
```

(continues on next page)

(continued from previous page)

```

integration.write-policy    functionality (querying the IAM database)
                             Integration testing: Tests the `write-policy`
                             function
test.lint                   Linting with `pylint` and `autopep8`
test.security               Runs `bandit` and `safety check`
unit.nose                   Unit testing: Runs unit tests using `nosetests`

# To run them, specify `invoke` plus the options:
invoke build.build-package

invoke integration.clean
invoke integration.initialize
invoke integration.analyze-policy
invoke integration.query
invoke integration.write-policy

invoke test.lint
invoke test.security

invoke unit.nose

```

13.3.3 Local Unit Testing and Integration Testing: Quick and Easy

We highly suggest that you run all the tests before pushing a significant commit. It would be painful to copy/paste all of those lines above - so we've compiled a test script in the *utils* folder.

Just run this from the root of the repository:

```
./utils/run_tests.sh
```

It will execute all of the tests that would normally be run during the TravisCI build. If you want to see if it will pass TravisCI, you can just run that quick command on your machine.

13.3.4 Running the Test Suite

We use *Nose* for unit testing. All tests are placed in the `tests` folder.

- Just run the following:

```
nosetests -v
```

- Alternatively, you can use *invoke*, as mentioned above:

```
invoke test.unit
```

Output:

```

test_overrides_yaml_config: Tests the format of the overrides yaml file for the RAM_
↪service ... ok
test_passing_overall_iam_action_override: Tests iam:CreateAccessKey ... ok
test_get_dependent_actions_double (test_actions.ActionsTestCase) ... ok
test_get_dependent_actions_several (test_actions.ActionsTestCase) ... ok
test_get_dependent_actions_single (test_actions.ActionsTestCase) ... ok
test_analyze_by_access_level: Test out calling this as a library ... ok

```

(continues on next page)

(continued from previous page)

```

test_get_actions_from_policy: Verify that the get_actions_from_policy function is_
↳grabbing the actions ... ok
test_get_actions_from_policy_file_with_explicit_actions: Verify that we can get a_
↳list of actions from a ... ok
test_get_actions_from_policy_file_with_wildcards: Verify that we can read the actions_
↳from a file, ... ok
test_remove_actions_not_matching_access_level: Verify remove_actions_not_matching_
↳access_level is working as expected ... ok
test_get_findings: Ensure that finding.get_findings() combines two risk findings for_
↳one policy properly. ... ok
test_get_findings_by_policy_name: Testing out the 'Findings' object ... ok
test_add_s3_permissions_management_arn (test_arn_action_group.ArnActionGroupTestCase)_
↳... ok
test_get_policy_elements (test_arn_action_group.ArnActionGroupTestCase) ... ok
test_update_actions_for_raw_arn_format (test_arn_action_group.ArnActionGroupTestCase)_
↳... ok
test_does_arn_match_case_1 (test_arns.ArnTestCase) ... ok
test_does_arn_match_case_2 (test_arns.ArnTestCase) ... ok
test_does_arn_match_case_4 (test_arns.ArnTestCase) ... ok
test_does_arn_match_case_5 (test_arns.ArnTestCase) ... ok
test_does_arn_match_case_6 (test_arns.ArnTestCase) ... ok
test_does_arn_match_case_bucket (test_arns.ArnTestCase) ... ok
test_determine_actions_to_expand: provide expanded list of actions, like ecr:* ... ok
test_minimize_statement_actions (test_minimize_wildcard_actions_
↳MinimizeWildcardActionsTestCase) ... ok
test_get_action_data: Tests function that gets details on a specific IAM Action. ..._
↳ok
test_get_actions_for_service: Tests function that gets a list of actions per AWS_
↳service. ... ok
test_get_actions_matching_condition_key: Tests a function that gathers all instances_
↳in ... ok
test_get_actions_that_support_wildcard_arns_only: Tests function that shows all ... ok
test_get_actions_with_access_level: Tests function that gets a list of actions in a ..
↳. ok
test_get_actions_with_arn_type_and_access_level: Tests a function that gets a list of_
↳... ok
test_get_arn_type_details: Tests function that grabs details about a specific ARN_
↳name ... ok
test_get_arn_types_for_service: Tests function that grabs arn_type and raw_arn pairs .
↳.. ok
test_get_condition_key_details: Tests function that grabs details about a specific_
↳condition key ... ok
test_get_condition_keys_for_service: Tests function that grabs a list of condition_
↳keys per service. ... ok
test_get_raw_arns_for_service: Tests function that grabs a list of raw ARNs per_
↳service ... ok
test_remove_actions_that_are_not_wildcard_arn_only: Tests function that removes_
↳actions from a list that ... ok
test_actions_template (test_template.TemplateTestCase) ... ok
test_crud_template (test_template.TemplateTestCase) ... ok
test_print_policy_with_actions_having_dependencies (test_write_policy_
↳WritePolicyActionsTestCase) ... ok
test_write_policy (test_write_policy.WritePolicyCrudTestCase) ... ok
Tests ARNs with the partition `aws-cn` instead of just `aws` ... ok
Tests ARNs with the partition `aws-us-gov` instead of `aws` ... ok
test_wildcard_when_not_necessary: Attempts bypass of CRUD mode wildcard-only ... ok
test_actions_missing_actions: write-policy actions if the actions block is missing ...
↳ ok

```

(continues on next page)

(continued from previous page)

```

test_allow_missing_access_level_categories_in_cfg: write-policy --crud when the YAML
↳file ... ok
test_allow_empty_access_level_categories_in_cfg: If the content of a list is an empty
↳string, it should sysexit ... ok
test_actions_missing_arn: write-policy actions command when YAML file block is
↳missing an ARN ... ok
test_actions_missing_description: write-policy when the YAML file is missing a
↳description ... ok
test_actions_missing_name: write-policy when the YAML file is missing a name ... ok

```

13.4 Project Structure

We'll focus mostly on the intent and approach of the major files (and subfolders) within the `policy_sentry/shared` directory:

13.4.1 Subfolders

Folders per command:

- The folders are mostly specific to their commands. For example, consider the files in the `policy_sentry/analysis` folder.
- **The files in this folder are specific to the *analyze* command**
 - They all can import from the *util* folder and the *shared* folder.
 - The files in this folder **don't import from other subfolders specific to other commands**, like *writing* or *downloading*. (Note: There is an occasional exception here of re-using functions from the *'querying'* folder)
 - Files in the *analysis* folder, to the *analyze* command. They don't import from each other, with the occasional exception of re-using functions from the *querying* folder. They all import common methods from the *util* folder and the *shared* folder as well.

Files:

- `shared/data/aws.sqlite3`: This is the pre-bundled IAM database. Third party packages can easily query the pre-bundled IAM database by connecting to the database like this: `db_session = connect_db('bundled')`
- `shared/data/audit/*.txt`: These text files are the pre-bundled audit files that you can use with the `analyze-iam-policy` command. Currently they are limited to privilege escalation and resource exposure. For more information, see the page on [Analyzing IAM Policies](#).
- `shared/data/docs/*.html`: These are HTML files wget'd from the [Actions, Resources, and Condition Keys](#) AWS documentation. This is used to build our database.
- `shared/data/access-level-overrides.yml`: This is created to override the access levels that AWS incorrectly states in their documentation. For instance, quite often, their service teams will say that an IAM action is "Tagging" when it really should be "Write" - for example, `secretsmanager:CreateSecret`.

13.4.2 Files and functions

TODO: Generate documentation automatically based on docstrings

13.5 Versioning

We try to follow [Semantic Versioning](#) as much as possible.

13.5.1 Version bumps

Just edit the `policy_sentry/bin/policy_sentry` file and update the `__version__` variable:

```
#!/usr/bin/env python
"""
    policy_sentry is a tool for generating least-privilege IAM Policies.
"""
__version__ = '0.6.3' # EDIT THIS
```

The `setup.py` file will automatically pick up the new version from that file for the package info. The `@click.version_option` decorator will also pick that up for the command line.

13.6 Roadmap

13.6.1 Condition Keys

Currently, Condition Keys are not supported by this script. For an example, see the KMS key Condition Key Table [here](#). Note: The database does create a table of condition keys in case we develop future support for it, but it isn't used yet.

13.6.2 Log-based policy generation

We are considering building functionality to:

- Use Amazon Athena to query CloudTrail logs from an S3 bucket for AWS IAM API calls, similar to [CloudTracker](#).
- **Instead of identifying the exact AWS IAM actions that were used, as CloudTracker currently does, we identify:**
 - Resource ARNs
 - Actions that indicate a CRUD level corresponding to that resource ARN. For example, if read access is granted to an S3 bucket folder path, assume all Read actions are needed for that folder path. Otherwise, we run into issues where CloudTrail actions and IAM actions don't match, which is a well documented issue by CloudTracker.
- **Query the logs to determine which principals touch which ARNs.**
 - For each IAM principal, create a list of ARNs.
 - For each ARN, plug that ARN into a `policy_sentry` yml file, and determine the CRUD level based on a lazy comparison of the action listed in the cloudtrail log vs the resource ARN.
 - And then run the `policy_sentry` yml file to generate an IAM policy that would have worked.

This was discussed in [the original Hacker News post](#)..

Policy Sentry can be used as a Python library. Check out this [documentation](#) for more information and examples.

14.1 Getting Started with the Library

When using Policy Sentry manually, you have to build a local database file with the initialize function.

However, if you are developing your own Python code and you want to import Policy Sentry as a third party package, you can skip the initialization and leverage the local database file that is bundled with the Python package itself.

This is especially useful for developers who wish to leverage Policy Sentry's capabilities that require the use of the IAM database (such as querying the IAM database table). This way, you don't have to initialize the database and can just query it immediately.

The code example is located [here](#). It is also shown below.

We've built a trick into the `connect_db` function that developers can specify to leverage the local database. The trick is to just use `'bundled'` as the single parameter for the `connect_db` method. See the example.

```
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_for_service

def example():
    db_session = connect_db('bundled') # This is the critical line. You just need to
    ↪ specify `bundled` as the parameter.
    actions = get_actions_for_service(db_session, 'cloud9') # Then you can leverage
    ↪ any method that requires access to the database.
    for action in actions:
        print(action)

if __name__ == '__main__':
    example()
```

Try running the code from the root of the repository:

```
./examples/library-usage/example.py
```

The results will look like this:

```
cloud9:createenvironmentec2
cloud9:createenvironmentmembership
cloud9:deleteenvironment
cloud9:deleteenvironmentmembership
cloud9:describeenvironmentmemberships
cloud9:describeenvironmentstatus
cloud9:describeenvironments
cloud9:getusersettings
cloud9:listenvironments
cloud9:updateenvironment
cloud9:updateenvironmentmembership
cloud9:updateusersettings
```

These are modules and functions that will be of interest for developers leveraging Policy Sentry as a Python library.

15.1 Querying

15.1.1 querying.all

IAM Database queries that are not specific to either the Actions, ARNs, or Condition Keys tables.

`policy_sentry.querying.all.get_all_actions(db_session)`

Gets a huge list of all IAM actions. This is used as part of the policyuniverse approach to minimizing IAM Policies to meet AWS-mandated character limits on policies.

Parameters `db_session` – SQLAlchemy database session object

Returns A list of all actions present in the database.

`policy_sentry.querying.all.get_all_service_prefixes(db_session)`

Gets all the AWS service prefixes from the actions table.

If the action table does NOT have specific IAM actions (and therefore only supports * actions), then it will not be included in the response.

Parameters `db_session` – The SQLAlchemy database session

Returns A list of all AWS service prefixes present in the table.

15.1.2 querying.actions

Methods that execute specific queries against the SQLite database for the ACTIONS table. This supports the policy_sentry query functionality

`policy_sentry.querying.actions.get_action_data(db_session, service, name)`

Get details about an IAM Action in JSON format.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*
- **name** – The name of an AWS IAM action, like *GetObject*.

Returns A dictionary containing metadata about an IAM Action.

`policy_sentry.querying.actions.get_actions_for_service(db_session, service)`

Get a list of available actions per AWS service

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*

Returns A list of actions

`policy_sentry.querying.actions.get_actions_matching_condition_key(db_session, service, condition_key)`

Get a list of actions under a service that allow the use of a specified condition key

Parameters

- **db_session** – SQLAlchemy database session
- **service** – A single AWS service prefix
- **condition_key** – The condition key to look for.

Returns A list of actions

`policy_sentry.querying.actions.get_actions_that_support_wildcard_arns_only(db_session, service)`

Get a list of actions that do not support restricting the action to resource ARNs.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – A single AWS service prefix, like *s3* or *kms*

Returns A list of actions

`policy_sentry.querying.actions.get_actions_with_access_level(db_session, service, access_level)`

Get a list of actions in a service under different access levels.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – A single AWS service prefix, like *s3* or *kms*
- **access_level** – An access level as it is written in the database, such as ‘Read’, ‘Write’, ‘List’, ‘Permissions Management’, or ‘Tagging’

Returns A list of actions

```
policy_sentry.querying.actions.get_actions_with_arn_type_and_access_level(db_session,
                                                                           ser-
                                                                           vice,
                                                                           re-
                                                                           source_type_name,
                                                                           ac-
                                                                           cess_level)
```

Get a list of actions in a service under different access levels, specific to an ARN format.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – A single AWS service prefix, like *s3* or *kms*
- **resource_type_name** – The ARN type name, like *bucket* or *key*

Returns A list of actions

```
policy_sentry.querying.actions.get_dependent_actions(db_session, actions_list)
```

Given a list of IAM Actions, query the database to determine if the action has dependent actions in the fifth column of the Resources, Actions, and Condition keys tables. If it does, add the dependent actions to the list, and return the updated list.

To get dependent actions for a single given IAM action, just provide the action as a list with one item, like this:
`get_dependent_actions(db_session, ['kms:CreateCustomKeystore'])`

Parameters

- **db_session** – SQLAlchemy database session object
- **actions_list** – A list of actions to use in querying the database for dependent actions

Returns Updated list of actions, including dependent actions if applicable.

```
policy_sentry.querying.actions.remove_actions_not_matching_access_level(db_session,
                                                                           ac-
                                                                           tions_list,
                                                                           ac-
                                                                           cess_level)
```

Given a list of actions, return a list of actions that match an access level

Parameters

- **db_session** – The SQLAlchemy database session
- **actions_list** – A list of actions
- **access_level** – ‘read’, ‘write’, ‘list’, ‘tagging’, or ‘permissions-management’

Returns Updated list of actions, where the actions not matching the requested access level are removed.

15.1.3 querying.arns

Methods that execute specific queries against the SQLite database for the ARN table. This supports the policy_sentry query functionality

```
policy_sentry.querying.arns.get_arn_type_details(db_session, service, name)
```

Get details about a resource ARN type name in JSON format. yo. stuff.

more.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*
- **name** – The name of a resource type, like *bucket* or *object*

Returns Metadata about an ARN type

`policy_sentry.querying.arns.get_arn_types_for_service(db_session, service)`
Get a list of available ARN short names per AWS service.

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*

Returns A list of ARN types, like *bucket* or *object*

`policy_sentry.querying.arns.get_raw_arns_for_service(db_session, service)`
Get a list of available raw ARNs per AWS service

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*

Returns A list of raw ARNs

15.1.4 querying.conditions

Methods that execute specific queries against the SQLite database for the CONDITIONS table. This supports the policy_sentry query functionality

`policy_sentry.querying.conditions.get_condition_key_details(db_session, service, condition_key_name)`

Get details about a specific condition key in JSON format

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *ec2* or *kms*
- **condition_key_name** – The name of a condition key, like *ec2:Vpc*

Returns Metadata about the condition key

`policy_sentry.querying.conditions.get_condition_keys_for_service(db_session, service)`

Get a list of available conditions per AWS service

Parameters

- **db_session** – SQLAlchemy database session object
- **service** – An AWS service prefix, like *s3* or *kms*

Returns A list of condition keys

15.2 Writing

15.2.1 command.write_policy

15.2.2 writing.policy

ArnActionGroup

class policy_sentry.writing.policy.ArnActionGroup

This class is critical to the creation of least privilege policies. It uses the SIDs as namespaces. The namespaces follow this format:

{Servicename}{Accesslevel}{Resourcetypeiname}

So, a resulting statement's SID might look like 'S3ListBucket'

add (*db_session, arn_list_from_user, access_level*)

This just adds the ARN, Service, and Access Level. ARN Format and Actions are not filled out. Example data can be found in the class ArnActionGroupTestCase in the testing folder.

Parameters

- **db_session** – SQLAlchemy database session
- **arn_list_from_user** – Just a list of resource ARNs.
- **access_level** – “Read”, “List”, “Tagging”, “Write”, or “Permissions management”

add_complete_entry (*arn_from_user, service, access_level, raw_arn_format, actions_list*)

Add a single entry with all the necessary fields filled out.

Parameters

- **arn_from_user** – The literal ARN that the user wants to put in the policy
- **service** – A service prefix, like s3
- **access_level** – Access level for the IAM action
- **raw_arn_format** – The raw arn that matches format of the arn_from_user
- **actions_list** – A list of actions

combine_policy_elements ()

Consolidate the policy elements by looking at where ARNs are used

does_action_exist (*action*)

Get boolean response for whether or not an action exists under any of the ARNs.

Parameters **action** – full action name, like s3:GetObject

Returns True or False

Return type bool

get_arns ()

Getter function for the ARNs object

Returns ARNs object

Return type dict

get_policy_elements (*db_session*)

Return the policy elements to the user.

Parameters `db_session` – database session.

Returns A dictionary that contains the SID namespace, a list of actions, and a list of resource ARNs that fall under this namespace

Return type dict

process_list_of_actions (*supplied_actions*, *db_session*)

Takes a list of actions, queries the database for corresponding arns, adds them to the object.

Parameters

- **supplied_actions** – A list of supplied actions
- **db_session** – SQLAlchemy database session object

Returns `arn_dict`: This is the compiled and updated dictionary after all necessary processing. This plugs into `create_policy`

Return type dict

process_resource_specific_acls (*cfg*, *db_session*)

Processes the YAML file for the resources per access level, and adds it to the object.

Parameters

- **cfg** – A dictionary loaded from the YAML file
- **db_session** – Database session

Returns the fully formed dict containing the ARNs and actions

Return type dict

remove_actions_duplicated_in_wildcard_resources ()

Removes actions from the object that are in a resource-specific ARN, as well as the `*` resource. For example, if `ssm:GetParameter` is restricted to a specific parameter path, as well as `*`, then we want to remove the `*` option to force least privilege.

remove_actions_not_matching_list (*actions_list*)

Parameters `actions_list` – List of actions to leave. All actions not in this list are removed

remove_sids_with_empty_action_lists ()

Now that we've removed a bunch of actions, if there are SID groups without any actions, remove them so we don't get SIDs with empty action lists

update_actions_for_raw_arn_format (*db_session*)

Considers the attribute values under each value in `self.arns`, and fills in the actions lists accordingly.

Parameters `db_session` – SQLAlchemy database session

`policy_sentry.writing.policy.create_policy_sid_namespace` (*service*, *access_level*, *resource_type_name*)

Simply generates the SID name. The SID groups ARN types that share an access level.

For example, S3 objects vs. SSM Parameter have different ARN types - as do S3 objects vs S3 buckets. That's how we choose to group them.

Parameters

- **service** – “ssm”
- **access_level** – “Read”
- **resource_type_name** – “parameter”

Returns SsmReadParameter

Return type str

`policy_sentry.writing.policy.remove_actions_that_are_not_wildcard_arn_only` (*db_session*,
actions_list)

Given a list of actions, remove the ones that CAN be restricted to ARNs, leaving only the ones that cannot.

Parameters

- **db_session** – SQL Alchemy database session object
- **actions_list** – A list of actions

Returns An updated list of actions

Return type list

15.2.3 writing.roles

Roles is a data holder for sublists. It supports the write-policy actions mode. In the future, this data structure should probably go away, or it should at least be renamed.

class `policy_sentry.writing.roles.Roles`

This is just a data holder for list of lists - used for running the YAML option. The sublists should all contain these: Indexes: 0 - Name 1 - Description 2 - ARN (string ARN of the role) 3 - Actions (list of AWS API Actions)

add_role (*role*)

Chain write-policy action mode files together

get_roles ()

Get the full data structure of action_list holder.

process_actions_config (*cfg*)

Given the YAML file used for the list of actions config, process it.

15.2.4 writing.template

Templates for the policy_sentry YAML files. These can be used for generating policies

`policy_sentry.writing.template.create_actions_template` (*name*)

Generate the Actions YAML template with Jinja2

`policy_sentry.writing.template.create_crud_template` (*name*)

Generate the CRUD YAML Template with Jinja2

15.2.5 writing.validate

Validation for the Policy Sentry YAML Templates.

`policy_sentry.writing.validate.check` (*conf_schema*, *conf*)

Validates a user-supplied JSON vs a defined schema. :param conf_schema: The Schema object that defines the required structure. :param conf: The user-supplied schema to validate against the required structure.

`policy_sentry.writing.validate.check_actions_schema` (*cfg*)

Determines whether the user-provided config matches the required schema for Actions mode

`policy_sentry.writing.validate.check_crud_schema(cfg)`

Determines whether the user-provided config matches the required schema for CRUD mode

15.2.6 writing.minimize

Functions for Minimizing statements, heavily borrowed from policyuniverse. <https://github.com/Netflix-Skunkworks/policyuniverse/>

IAM Policies have character limits, which apply to individual policies, and there are also limits on the total aggregate policy sizes. As such, it is not possible to use exhaustive list of explicit IAM actions. To have granular control of specific IAM policies, we must use wildcards on IAM Actions, only in a programmatic manner.

This is typically performed by humans by reducing policies to `s3:Get*`, `ec2:Describe*`, and other approaches of the sort.

Netflix's PolicyUniverse has address

<https://aws.amazon.com/iam/faqs/> Q: How many policies can I attach to an IAM role? * For inline policies: You can add as many inline policies as you want to a user, role, or group, but

the total aggregate policy size (the sum size of all inline policies) per entity cannot exceed the following limits: - User policy size cannot exceed 2,048 characters. - Role policy size cannot exceed 10,240 characters. - Group policy size cannot exceed 5,120 characters.

- For managed policies: You can add up to 10 managed policies to a user, role, or group.
- The size of each managed policy cannot exceed 6,144 characters.

`policy_sentry.writing.minimize.check_min_permission_length(permission, min-
chars=None)`

Adapted version of policyuniverse's `_check_permission_length`. We are commenting out the skipping prefix message https://github.com/Netflix-Skunkworks/policyuniverse/blob/master/policyuniverse/expander_minimizer.py#L111

`policy_sentry.writing.minimize.get_denied_prefixes_from_desired(desired_actions, all_actions)`

Adapted version of policyuniverse's `_get_denied_prefixes_from_desired`, here: https://github.com/Netflix-Skunkworks/policyuniverse/blob/master/policyuniverse/expander_minimizer.py#L101

`policy_sentry.writing.minimize.minimize_statement_actions(desired_actions, all_actions, min-
chars=None)`

This is a condensed version of policyuniverse's `minimize_statement_actions`, changed for our purposes. https://github.com/Netflix-Skunkworks/policyuniverse/blob/master/policyuniverse/expander_minimizer.py#L123

15.3 Analyzing

15.3.1 analysis.analyze

Functions to support the analyze capability in this tool

`policy_sentry.analysis.analyze.analyze_by_access_level(db_session, policy_json, access_level)`

Determine if a policy has any actions with a given access level. This is particularly useful when determining who has 'Permissions management' level access

Parameters

- **db_session** – SQLAlchemy database session

- **policy_json** – a dictionary representing the AWS JSON policy
- **access_level** – The normalized access level - either 'read', 'list', 'write', 'tagging', or 'permissions-management'

```
policy_sentry.analysis.analyze.analyze_policy_directory(db_session,          pol-
                                                         icy_directory,          ac-
                                                         count_id, from_audit_file,
                                                         finding_type,          ex-
                                                         cluded_role_patterns)
```

Audits a directory of policy JSON files.

Parameters

- **db_session** – SQLAlchemy database session object
- **policy_directory** – The file directory where the policies are stored
- **account_id** – The AWS Account ID
- **from_audit_file** – The file containing the list of problematic actions
- **finding_type** – The type of finding - resource_exposure, privilege_escalation, network_exposure, or credentials_exposure

Returns A dictionary of findings with the policy names as keys.

```
policy_sentry.analysis.analyze.analyze_policy_file(db_session,  policy_file,  ac-
                                                         count_id, from_audit_file, find-
                                                         ing_type, excluded_role_patterns)
```

Given a policy file, determine risky actions based on a separate file containing a list of actions. If it matches a policy exclusion pattern from the report-config.yml file, that policy file will be skipped.

Parameters

- **db_session** – SQLAlchemy database session object
- **policy_file** – The path to the policy file to be evaluated
- **account_id** – The AWS Account ID
- **from_audit_file** – The file containing the list of problematic actions
- **finding_type** – The type of finding - resource_exposure, privilege_escalation, network_exposure, or credentials_exposure
- **excluded_role_patterns** – A RegEx pattern for excluding policy names from evaluation.

Returns False if the policy name matches excluded role patterns, or if it does not, a dictionary containing the findings.

Return type dict

```
policy_sentry.analysis.analyze.determine_actions_to_expand(db_session,          ac-
                                                         tion_list)
```

Determine if an action needs to get expanded from its wildcard

Parameters

- **db_session** – A SQLAlchemy database session object
- **action_list** – A list of actions

Returns A list of actions

Return type list

`policy_sentry.analysis.analyze.determine_risky_actions` (*requested_actions*, *audit_file*)
compare the actions in the policy against the audit file of high risk actions

Parameters

- **requested_actions** – A list of the actions that are requested by the policy under evaluation
- **audit_file** – The absolute path to the file that contains a list of IAM action to evaluate.

Returns a list of any actions that are included in the file of risky actions

`policy_sentry.analysis.analyze.expand` (*action*, *db_session*)
expand the action wildcards into a full action

Parameters

- **action** – An action in the form with a wildcard - like s3:Get*, or s3:L*
- **db_session** – SQLAlchemy database session object

Returns A list of all the expanded actions (like actions matching s3:Get*)

Return type list

`policy_sentry.analysis.analyze.read_risky_iam_permissions_text_file` (*audit_file*)
read in the audit file of high risk actions

Parameters **audit_file** – Path to the file containing a list of risky actions

Return risky_actions A list of actions from the file

15.4 Utilities

15.4.1 util.policy_files

A few methods for parsing policies.

`policy_sentry.util.policy_files.get_actions_from_json_policy_file` (*file*)
read the json policy file and return a list of actions

`policy_sentry.util.policy_files.get_actions_from_policy` (*data*)
Given a policy dictionary, create a list of the actions

15.4.2 util.arns

Functions to use for parsing ARNs, matching ARN types, and getting the right fragment/component from an ARN string,

`policy_sentry.util.arns.arn_has_colons` (*arn*)
Given an ARN, determine if the ARN has colons in it. Just useful for the hacky methods for parsing ARN namespaces. See <http://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html> for more details on ARN namespaces.

`policy_sentry.util.arns.arn_has_slash` (*arn*)
Given an ARN, determine if the ARN has a slash in it. Just useful for the hacky methods for parsing ARN namespaces. See <http://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html> for more details on ARN namespaces.

`policy_sentry.util.arns.does_arn_match(arn_to_test, arn_in_database)`

Given two ARNs, determine if they match. The cases supported are outlined below.

Case 1: `arn:partition:service:region:account-id:resource` Case 2: `arn:partition:service:region:account-id:resourcetype/resource` Case 3: `arn:partition:service:region:account-id:resourcetype/resource/qualifier`
 Case 4: `arn:partition:service:region:account-id:resourcetype/resource:qualifier` Case 5: `arn:partition:service:region:account-id:resourcetype:resource` Case 6: `arn:partition:service:region:account-id:resourcetype:resource:qualifier` Source: <https://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html#genref-arns>

Parameters `arn` – ARN to parse

Returns result of whether or not the ARNs match

`policy_sentry.util.arns.get_account_from_arn(arn)`

Given an ARN, return the account ID in the ARN, if it is available. In certain cases like S3 it is not

`policy_sentry.util.arns.get_partition_from_arn(arn)`

Given an ARN string, return the partition string. This is usually `aws` unless you are in C2S or AWS GovCloud.

`policy_sentry.util.arns.get_region_from_arn(arn)`

Given an ARN, return the region in the ARN, if it is available. In certain cases like S3 it is not

`policy_sentry.util.arns.get_resource_from_arn(arn)`

Given an ARN, parse it according to ARN namespacing and return the resource. See <http://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html> for more details on ARN namespacing.

`policy_sentry.util.arns.get_resource_path_from_arn(arn)`

Given an ARN, parse it according to ARN namespacing and return the resource path. See <http://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html> for more details on ARN namespacing.

`policy_sentry.util.arns.get_service_from_arn(arn)`

Given an ARN string, return the service

`policy_sentry.util.arns.parse_arn(arn)`

Given an ARN, split up the ARN into the ARN namespacing schema dictated by the AWS docs.

15.4.3 util.file

Functions that relate to manipulating files, loading files, and managing filepaths.

`policy_sentry.util.file.check_valid_file_path(file)`

Checks if the file path is valid.

Parameters `file` – The file to check.

Returns True if it exists, False if it does not

Return type bool

`policy_sentry.util.file.create_directory_if_it_doesnt_exist(directory)`

Equivalent of `mkdir -p`

`policy_sentry.util.file.list_files_in_directory(directory)`

Equivalent of `ls` command, and return the list of files

`policy_sentry.util.file.read_this_file(filename)`

Read a file at a path and return the lines from each file

`policy_sentry.util.file.read_yaml_file(filename)`

Reads a YAML file, safe loads, and returns the dictionary

Parameters `filename` – name of the yaml file

Returns dictionary of YAML file contents

:

`policy_sentry.util.file.write_json_file(filename, json_contents)`

Description: Writes a YAML file :param json_contents: a dictionary used to build the JSON. This is the IAM Policy built by write_policy functions. :param filename: name of the yaml file, which should include the path

15.4.4 util.actions

Text operations specific to IAM actions

`policy_sentry.util.actions.get_action_name_from_action(action)`

Returns the lowercase action name from a `service:action` combination :param action: `ec2:DescribeInstance` :return: `describeinstance`

`policy_sentry.util.actions.get_full_action_name(service, action_name)`

Gets the proper formatting for an action - the service, plus colon, plus action name. :param service: service name, like `s3` :param action_name: action name, like `createbucket` :return: the resulting string

`policy_sentry.util.actions.get_lowercase_action_list(action_list)`

Given a list of actions, return the list but in lowercase format

`policy_sentry.util.actions.get_service_from_action(action)`

Returns the service name from a `service:action` combination :param action: `ec2:DescribeInstance` :return: `ec2`

Examples

These are examples for the modules and functions that will be of interest for developers leveraging Policy Sentry as a Python library.

16.1 Querying

The following are examples of how to leverage some of the functions available from Policy Sentry. The functions selected are likely to be of most interest to other developers.

16.1.1 All

querying.all.get_all_services

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.all import get_all_service_prefixes

if __name__ == '__main__':
    db_session = connect_db('bundled')
    all_service_prefixes = get_all_service_prefixes(db_session)
    print(all_service_prefixes)

"""
Output:

A list of every service prefix (like 'kms' or 's3') available in the IAM database.
Note that this will not include services that do not support any ARN types, like AWS_
↳ IQ.
"""
```

querying.all.get_all_actions

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.all import get_all_actions

if __name__ == '__main__':
    db_session = connect_db('bundled')
    all_actions = get_all_actions(db_session)
    print(all_actions)

"""
Output:

Every IAM action available across all services, without duplicates
"""
```

16.1.2 Actions

querying.actions.get_action_data

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_action_data
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_action_data(db_session, 'ram', 'createresourceshare')
    print(json.dumps(output, indent=4))

"""
Output:

{
  'ram': [
    {
      'action': 'ram:createresourceshare',
      'description': 'Create resource share with provided resource(s) and/or_
↪principal(s) ',
      'access_level': 'Permissions management',
      'resource_arn_format': 'arn:${Partition}:ram:${Region}:${Account}
↪:resource-share/${ResourcePath}',
      'condition_keys': [
        'ram:RequestedResourceType',
        'ram:ResourceArn',
        'ram:RequestedAllowsExternalPrincipals'
      ],
      'dependent_actions': None
    },
    {
      'action': 'ram:createresourceshare',
      'description': 'Create resource share with provided resource(s) and/or_
↪principal(s) ',
```

(continues on next page)

(continued from previous page)

```

        'access_level': 'Permissions management',
        'resource_arn_format': '*',
        'condition_keys': [
            'aws:RequestTag/${TagKey}',
            'aws:TagKeys'
        ],
        'dependent_actions': None
    }
]
}
"""

```

querying.actions.get_actions_for_service

```

#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_for_service
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_actions_for_service(db_session, 'cloud9')
    print(json.dumps(output, indent=4))

"""
Output:

[
    'ram:acceptresourceshareinvitation',
    'ram:associateresourceshare',
    'ram:createresourceshare',
    'ram:deleteresourceshare',
    'ram:disassociateresourceshare',
    'ram:enablessharingwithawsorganization',
    'ram:rejectresourceshareinvitation',
    'ram:updateresourceshare'
]
"""

```

querying.actions.get_actions_matching_condition_key

```

#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_matching_condition_key
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_actions_matching_condition_key(db_session, "ses",
↪ "ses:FeedbackAddress")
    print(json.dumps(output, indent=4))

"""

```

(continues on next page)

(continued from previous page)

Output:

```
[
    'ses:sendemail',
    'ses:sendbulktemplatedemail',
    'ses:sendcustomverificationemail',
    'ses:sendemail',
    'ses:sendrawemail',
    'ses:sendtemplatedemail'
]
"""
```

querying.actions.get_actions_supporting_wilcards_only

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_matching_condition_key
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_actions_matching_condition_key(db_session, "ses",
↪ "ses:FeedbackAddress")
    print(json.dumps(output, indent=4))

"""
Output:

[
    'ses:sendemail',
    'ses:sendbulktemplatedemail',
    'ses:sendcustomverificationemail',
    'ses:sendemail',
    'ses:sendrawemail',
    'ses:sendtemplatedemail'
]
"""
```

querying.actions.get_actions_with_access_levels

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_with_access_level
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_actions_with_access_level(db_session, 's3', 'Permissions management')
    print(json.dumps(output, indent=4))

"""
Output:
```

(continues on next page)

(continued from previous page)

```

s3:bypassgovernanceretention
s3:deleteaccesspointpolicy
s3:deletebucketpolicy
s3:objectowneroverridetobucketowner
s3:putaccesspointpolicy
s3:putaccountpublicaccessblock
s3:putbucketacl
s3:putbucketpolicy
s3:putbucketpublicaccessblock
s3:putobjectacl
s3:putobjectversionacl
"""

```

querying.actions.get_actions_with_arn_type_and_access_level

```

#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_actions_with_arn_type_and_access_level
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_actions_with_arn_type_and_access_level(db_session, "ram", "resource-
↪share", "Permissions management")
    print(json.dumps(output, indent=4))

"""
Output:

[
    'ram:associateresourceshare',
    'ram:createresourceshare',
    'ram:deleteresourceshare',
    'ram:disassociateresourceshare',
    'ram:updateresourceshare'
]
"""

```

querying.actions.get_dependent_actions

```

#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.actions import get_dependent_actions
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_dependent_actions(db_session, ["ec2:associateiaminstanceprofile"])
    print(json.dumps(output, indent=4))

"""
Output:

```

(continues on next page)

(continued from previous page)

```
[
    "ec2:associateiaminstanceprofile",
    "iam:passrole"
]
"""
```

16.1.3 ARNs

querying.arns.get_arn_type_details

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.arns import get_arn_type_details
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_arn_type_details(db_session, "cloud9", "environment")
    print(json.dumps(output, indent=4))

"""
Output:

{
    "resource_type_name": "environment",
    "raw_arn": "arn:${Partition}:cloud9:${Region}:${Account}:environment:${ResourceId}"
    ↪,
    "condition_keys": None
}
"""
```

querying.arns.get_arn_types_for_service

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.arns import get_arn_types_for_service
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_arn_types_for_service(db_session, "s3")
    print(json.dumps(output, indent=4))

"""
Output:

{
    "accesspoint": "arn:${Partition}:s3:${Region}:${Account}:accesspoint/$"
    ↪{AccessPointName}",
    "bucket": "arn:${Partition}:s3:::${BucketName}",
    "object": "arn:${Partition}:s3:::${BucketName}/${ObjectName}",
    "job": "arn:${Partition}:s3:${Region}:${Account}:job/${JobId}",

```

(continues on next page)

(continued from previous page)

```
}
"""
```

querying.arns.get_raw_arns_for_service

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.arns import get_raw_arns_for_service
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_raw_arns_for_service(db_session, "s3")
    print(json.dumps(output, indent=4))

"""
Output:

[
    "arn:${Partition}:s3:${Region}:${Account}:accesspoint/${AccessPointName}",
    "arn:${Partition}:s3:::${BucketName}",
    "arn:${Partition}:s3:::${BucketName}/${ObjectName}",
    "arn:${Partition}:s3:${Region}:${Account}:job/${JobId}"
]
"""
```

16.1.4 Conditions

querying.conditions.get_condition_key_details

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.conditions import get_condition_key_details
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_condition_key_details(db_session, "cloud9", "cloud9:Permissions")
    print(json.dumps(output, indent=4))

"""
Output:

{
    "name": "cloud9:Permissions",
    "description": "Filters access by the type of AWS Cloud9 permissions",
    "condition_value_type": "string"
}
"""
```

querying.conditions.get_condition_keys_for_service

```
#!/usr/bin/env python
from policy_sentry.shared.database import connect_db
from policy_sentry.querying.conditions import get_condition_keys_for_service
import json

if __name__ == '__main__':
    db_session = connect_db('bundled')
    output = get_condition_keys_for_service(db_session, "cloud9")
    print(json.dumps(output, indent=4))

"""
Output:

[
    'cloud9:EnvironmentId',
    'cloud9:EnvironmentName',
    'cloud9:InstanceType',
    'cloud9:Permissions',
    'cloud9:SubnetId',
    'cloud9:UserArn'
]
"""
```

Implementation Strategy

In the context of your overall organization strategy for AWS IAM, we recommend using a few measures for locking down your AWS environments with IAM:

1. Use [policy_sentry](#) to create [Identity-based policies](#)
2. Use Service Control Policies (SCPs) to lock down available API calls per account.
 - A great collection of SCPs can be found [on asecure.cloud](#).
 - Control Tower has some excellent guidance on strategy for SCPs in their documentation. Note that they call it “Guardrails” but they are mostly SCPs. See the docs [here](#)
3. Use [Repokid](#) to revoke out of date policies as your application/roles mature.
4. Use [Resource-based policies](#) for all services that support them.
 - A list of which services support resource-based policies can be found [in the AWS documentation here](#).
5. Never provision infrastructure manually; use Infrastructure as Code
 - I highly suggest Terraform for IAC over other alternatives such as CloudFormation, Chef, or Puppet. Yevgeniy Brikman explains the reasons very well in [this Gruntwork.io blog post](#).
 - I also suggest reading HashiCorp’s [Unlocking the Cloud Operating Model Whitepaper](#).

This document covers:

- Elements of an IAM Policy
- Breakdown of the tables for Actions, Resources, and Condition keys per service
- Generally how policy_sentry uses these tables to generate IAM Policies

18.1 IAM Policy Elements

The following IAM JSON Policy elements are included in policy_sentry-generated IAM Policies:

- **Version:** specifies policy language versions dictated by AWS. There are two options - 2012-10-17 and 2008-10-17. policy_sentry generates policies for the most recent policy language - 2012-10-17
- **Statement:** There is one **statement array** per policy, with multiple statements/SIDs inside that statement. The elements of a single statement/SID are listed below.
 - **SID:** Statement ID. Optional identifier for the policy statement. SID values can be assigned to each statement in a statement array.
 - **Effect:** Allow or explicit Deny. If there is any overlap on an action or actions with Allow vs. Deny, the Deny effect overrides the Allow.
 - **Action:** This refers to the IAM action - i.e., s3:GetObject, or ec2:DescribeInstances. Action text in a statement can have wildcards included: for example, ec2:* covers all EC2 actions, and ec2:Describe* covers all EC2 actions prefixed with Describe - such as DescribeInstances, DescribeInstanceAttributes, etc.
 - **Resource:** This refers to an Amazon Resource Name (ARN) that the Action can be performed against. There are differences in ARN format per service. Those differences can be viewed [in the AWS Docs on ARNs and Namespaces](#)

The ones we don't use in this tool:

- **Condition** (will be added in a future release)

- Principal
- NotPrincipal
- NotResource

18.2 Actions, Resources, and Condition Keys Per Service

If you *ever* write or review IAM Policies, you should bookmark the documentation page for AWS Actions, Resources, and Context Keys [here](#)

This documentation is the seed source for the database that we create in policy_sentry. It contains tables for **(1) Actions**, **(2) Resources/ARNs**, and **(3) Condition Keys** for each service. This documentation is of critical importance because **every IAM action for every IAM service has different ARNs that it can apply to, and different Condition Keys that it can apply to.**

18.2.1 Action Table

Consider the Action table snippet from KMS shown below (source documentation can be viewed [on the KMS documentation here](#)).

Actions	Access Level	Resource Types	Condition Keys	Dependent Actions
kms:CreateGrant	Permissions management	key*		
		•	kms:CallerAccount	
kms:CreateCustomKeyStore	Write	•		cloudhsm:DescribeClusters

As you can see, the Actions Table contains these columns:

- **Actions:** The name of the IAM Action
- **Access Level:** how the action is classified. This is limited to List, Read, Write, Permissions management, or Tagging.
 - This classification can help you understand the level of access that an action grants when you use it in a policy.
 - For more information about access levels, see [Understanding Access Level Summaries Within Policy Summaries](#).
- **Condition Keys:** The condition key available for that action. There are some service specific ones that will contain the service namespace (i.e., `ec2`, or in this case, `kms`). Sometimes, there are AWS-level condition keys that are available to only some actions within some services, such as `aws:SourceAccount`. If those are available to the action, they will be supplied in that column.
- **Dependent Actions:** Some actions require that other actions can be executed by the IAM Principal. The example above indicates that in order to call `kms:CreateCustomKeyStore`, you must be able to also execute `cloudhsm:DescribeClusters`.

And most importantly to the context of this tool, there is the Resource Types column:

- **Resource Types:** This indicates whether the action supports resource-level permissions - i.e., *restricting IAM Actions by ARN*. If there is a value here, it points to the ARN Table shown later in the documentation.

- In the example above, you can see that `kms:CreateCustomKeyStore`'s Resource Types cell is blank; this indicates that `kms:CreateCustomKeyStore` can **only** have `*` as the resource type.
- Conversely, for `kms:CreateGrant`, the action can have either (1) `*` as the resource type, or `key*` as the resource type. The ARN format is not actually `key*`, it just points to that ARN format in the ARN Table explained below.

18.2.2 ARN Table

Consider the KMS ARN Table shown below (the source documentation can be viewed [on the AWS website here](#)).

Resource Types	ARN	Condition Keys
alias	<code>arn:\${Partition}:kms:\${Region}:\${Account}:alias/\${Alias}</code>	
key	<code>arn:\${Partition}:kms:\${Region}:\${Account}:key/\${KeyId}</code>	

The ARN Table has three fields:

- **Resource Types:** The name of the resource type. This corresponds to the “Resource Types” field in the Action table. In the example above, the types are:
 - `alias`
 - `key`
- **ARN:** This shows the required ARN format that can be specified in IAM policies for the IAM Actions that allow this ARN format. In the example above the ARN types are:
 - `arn:${Partition}:kms:${Region}:${Account}:alias/${Alias}`
 - `arn:${Partition}:kms:${Region}:${Account}:key/${KeyId}`
- **Condition Keys:** This specifies condition context keys that you can include in an IAM policy statement only when both (1) this resource and (2) a supporting action from the table above are included in the statement.

18.2.3 Condition Keys Table

There is also a Condition Keys table. An example is shown below.

Condition Keys	Type	Description
<code>kms:BypassPolicyLockoutSafetyCheck</code>	Boolean	Controls access to the <code>CreateKey</code> and <code>PutKeyPolicy</code> operations based on the value of the <code>BypassPolicyLockoutSafetyCheck</code> parameter in the request.
<code>kms:CallerAccount</code>	String	Controls access to specified AWS KMS operations based on the AWS account ID of the caller. You can use this condition key to allow or deny access to all IAM users and roles in an AWS account in a single policy statement.

Note: While `policy_sentry` does import the Condition Keys table into the database, it does not currently provide functionality to insert these condition keys into the policies. This is due to the complexity of each condition key, and the dubious viability of mandating those condition keys for every IAM policy.

We might support the Global Condition keys for IAM policies in the future, perhaps to be supplied via a user config file, but that functionality is not on the roadmap at this time. For more information on Global Condition Keys, see [this documentation](#).

18.2.4 References

- [ARN Formats and Service Namespaces](#)
- [IAM Policy Elements](#)
- [IAM Actions, Resources, and Context Keys per service](#)
- [Actions Table explanation](#)
- [ARN Table explanation](#)
- [Condition Keys Table explanation](#)
- [Global Condition Keys](#)

Minimization

This document explains the approach in the file titled `policy_sentry/shared/minimize.py`, which is heavily borrowed from Netflix's [policyuniverse](#)

IAM Policies have character limits, which apply to individual policies, and there are also limits on the total aggregate policy sizes. As such, it is not possible to use exhaustive list of explicit IAM actions. To have granular control of specific IAM policies, we must use wildcards on IAM Actions, only in a programmatic manner.

This is typically performed by humans by reducing policies to `s3:Get*`, `ec2:Describe*`, and other approaches of the sort.

Netflix's `PolicyUniverse1` has addressed this problem using a few functions that we borrowed directly, and slightly modified. All of these functions are inside the aforementioned `minimize.py` file, and are also listed below:

- `get_denied_prefixes_from_desired`
- `check_min_permission_length`
- `minimize_statement_actions`

We modified the functions, in short, because of how we source our list of IAM actions. `Policyuniverse` leverages a file titled `data.json`, which appears to be a manually altered version of the `policies.js` file included as part of the [AWS Policy Generator website](#). However, that page is not updated as frequently. It also does not include the same details that we get from the [Actions, Resources, and Condition Keys](#) page, like the Dependent Actions Field, service-specific conditions, and most importantly the multiple ARN format types that can apply to any particular IAM Action.

See the AWS IAM FAQ page for supporting details on IAM Size. For your convenience, the relevant text is clipped below.

Q: How many policies can I attach to an IAM role?

- For inline policies: You can add as many inline policies as you want to a user, role, or group, but the total aggregate policy size (the sum size of all inline policies) per entity cannot exceed the following limits:
 - User policy size cannot exceed 2,048 characters.
 - Role policy size cannot exceed 10,240 characters.
 - Group policy size cannot exceed 5,120 characters.

- For managed policies: You can add up to 10 managed policies to a user, role, or group.
- The size of each managed policy cannot exceed 6,144 characters.

CHAPTER 20

Indices and tables

- `modindex`

p

- `policy_sentry.analysis.analyze`, 62
- `policy_sentry.querying.actions`, 55
- `policy_sentry.querying.all`, 55
- `policy_sentry.querying.arns`, 57
- `policy_sentry.querying.conditions`, 58
- `policy_sentry.util.actions`, 66
- `policy_sentry.util.arns`, 64
- `policy_sentry.util.file`, 65
- `policy_sentry.util.policy_files`, 64
- `policy_sentry.writing.minimize`, 62
- `policy_sentry.writing.policy`, 59
- `policy_sentry.writing.roles`, 61
- `policy_sentry.writing.template`, 61
- `policy_sentry.writing.validate`, 61

A

`add()` (*policy_sentry.writing.policy.ArnActionGroup* method), 59
`add_complete_entry()` (*policy_sentry.writing.policy.ArnActionGroup* method), 59
`add_role()` (*policy_sentry.writing.roles.Roles* method), 61
`analyze_by_access_level()` (in module *policy_sentry.analysis.analyze*), 62
`analyze_policy_directory()` (in module *policy_sentry.analysis.analyze*), 63
`analyze_policy_file()` (in module *policy_sentry.analysis.analyze*), 63
`arn_has_colons()` (in module *policy_sentry.util.arns*), 64
`arn_has_slash()` (in module *policy_sentry.util.arns*), 64
`ArnActionGroup` (class in *policy_sentry.writing.policy*), 59

C

`check()` (in module *policy_sentry.writing.validate*), 61
`check_actions_schema()` (in module *policy_sentry.writing.validate*), 61
`check_crud_schema()` (in module *policy_sentry.writing.validate*), 61
`check_min_permission_length()` (in module *policy_sentry.writing.minimize*), 62
`check_valid_file_path()` (in module *policy_sentry.util.file*), 65
`combine_policy_elements()` (*policy_sentry.writing.policy.ArnActionGroup* method), 59
`create_actions_template()` (in module *policy_sentry.writing.template*), 61
`create_crud_template()` (in module *policy_sentry.writing.template*), 61
`create_directory_if_it_doesnt_exist()`

(in module *policy_sentry.util.file*), 65

`create_policy_sid_namespace()` (in module *policy_sentry.writing.policy*), 60

D

`determine_actions_to_expand()` (in module *policy_sentry.analysis.analyze*), 63
`determine_risky_actions()` (in module *policy_sentry.analysis.analyze*), 63
`does_action_exist()` (*policy_sentry.writing.policy.ArnActionGroup* method), 59
`does_arn_match()` (in module *policy_sentry.util.arns*), 64

E

`expand()` (in module *policy_sentry.analysis.analyze*), 64

G

`get_account_from_arn()` (in module *policy_sentry.util.arns*), 65
`get_action_data()` (in module *policy_sentry.querying.actions*), 55
`get_action_name_from_action()` (in module *policy_sentry.util.actions*), 66
`get_actions_for_service()` (in module *policy_sentry.querying.actions*), 56
`get_actions_from_json_policy_file()` (in module *policy_sentry.util.policy_files*), 64
`get_actions_from_policy()` (in module *policy_sentry.util.policy_files*), 64
`get_actions_matching_condition_key()` (in module *policy_sentry.querying.actions*), 56
`get_actions_that_support_wildcard_arns_only()` (in module *policy_sentry.querying.actions*), 56
`get_actions_with_access_level()` (in module *policy_sentry.querying.actions*), 56
`get_actions_with_arn_type_and_access_level()` (in module *policy_sentry.querying.actions*), 56

`get_all_actions()` (in module `policy_sentry.querying.all`), 55
`get_all_service_prefixes()` (in module `policy_sentry.querying.all`), 55
`get_arn_type_details()` (in module `policy_sentry.querying.arns`), 57
`get_arn_types_for_service()` (in module `policy_sentry.querying.arns`), 58
`get_arns()` (`policy_sentry.writing.policy.ArnActionGroup` method), 59
`get_condition_key_details()` (in module `policy_sentry.querying.conditions`), 58
`get_condition_keys_for_service()` (in module `policy_sentry.querying.conditions`), 58
`get_denied_prefixes_from_desired()` (in module `policy_sentry.writing.minimize`), 62
`get_dependent_actions()` (in module `policy_sentry.querying.actions`), 57
`get_full_action_name()` (in module `policy_sentry.util.actions`), 66
`get_lowercase_action_list()` (in module `policy_sentry.util.actions`), 66
`get_partition_from_arn()` (in module `policy_sentry.util.arns`), 65
`get_policy_elements()` (`policy_sentry.writing.policy.ArnActionGroup` method), 59
`get_raw_arns_for_service()` (in module `policy_sentry.querying.arns`), 58
`get_region_from_arn()` (in module `policy_sentry.util.arns`), 65
`get_resource_from_arn()` (in module `policy_sentry.util.arns`), 65
`get_resource_path_from_arn()` (in module `policy_sentry.util.arns`), 65
`get_roles()` (`policy_sentry.writing.roles.Roles` method), 61
`get_service_from_action()` (in module `policy_sentry.util.actions`), 66
`get_service_from_arn()` (in module `policy_sentry.util.arns`), 65

L

`list_files_in_directory()` (in module `policy_sentry.util.file`), 65

M

`minimize_statement_actions()` (in module `policy_sentry.writing.minimize`), 62

P

`parse_arn()` (in module `policy_sentry.util.arns`), 65
`policy_sentry.analysis.analyze` (module), 62
`policy_sentry.querying.actions` (module), 55
`policy_sentry.querying.all` (module), 55
`policy_sentry.querying.arns` (module), 57
`policy_sentry.querying.conditions` (module), 58
`policy_sentry.util.actions` (module), 66
`policy_sentry.util.arns` (module), 64
`policy_sentry.util.file` (module), 65
`policy_sentry.util.policy_files` (module), 64
`policy_sentry.writing.minimize` (module), 62
`policy_sentry.writing.policy` (module), 59
`policy_sentry.writing.roles` (module), 61
`policy_sentry.writing.template` (module), 61
`policy_sentry.writing.validate` (module), 61
`process_actions_config()` (`policy_sentry.writing.roles.Roles` method), 61
`process_list_of_actions()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60
`process_resource_specific_acls()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60

R

`read_risky_iam_permissions_text_file()` (in module `policy_sentry.analysis.analyze`), 64
`read_this_file()` (in module `policy_sentry.util.file`), 65
`read_yaml_file()` (in module `policy_sentry.util.file`), 65
`remove_actions_duplicated_in_wildcard_resources()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60
`remove_actions_not_matching_access_level()` (in module `policy_sentry.querying.actions`), 57
`remove_actions_not_matching_list()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60
`remove_actions_that_are_not_wildcard_arn_only()` (in module `policy_sentry.writing.policy`), 61
`remove_sids_with_empty_action_lists()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60
`Roles` (class in `policy_sentry.writing.roles`), 61

U

`update_actions_for_raw_arn_format()` (`policy_sentry.writing.policy.ArnActionGroup` method), 60

W

`write_json_file()` (*in module policy_{sentry}.util.file*), [66](#)