

Sourced Group

Detective Controls, Python and AWS

1st November 2019

Contents

01

Introduction to Sourced

02

Cloud Overview

03

Detective and Preventative
Controls

04

Problem Definition

05

Selecting the Right Tool

About Me

- Consultant at Sourced Group
- ❤️ Python
- Grew up near Sydney, Australia
- Has never been on Bondi Rescue
- Volunteer in Rescue with NSW State Emergency Service



About Sourced

- Sourced Group was founded in 2010 by a team of consultants with deep financial services experience
- 150+ people across 3 regions
 - Sydney, Australia
 - Toronto, Canada
 - Singapore
- Deep specialisation in architecture and delivery on Amazon Web Services
- Experience with clients in a number of industries:
 - Financial Services
 - Aviation
 - Telco
 - Health
 - Media



Premier
Consulting
Partner

Migration Competency

DevOps Competency

Channel Partner

Public Sector Partner

Some of our Clients

Proven track record of delivering transformational projects in the enterprise



CommonwealthBank



SINGAPORE AIRLINES



mastercard.



QANTAS



Python 101

What can you use Python for?

- Web applications – Django, Flask
- Deep Learning & Machine learning – Pytorch, Tensorflow
- “Glue” scripts to connect systems & applications
- General programming – ships natively on many systems

```
1  #!/usr/bin/env python3
2
3  print("Hello world!")
```

PYTORCH



TensorFlow

What is the Cloud?

What is the Cloud?

- You used to have to build or rent space in a datacenter
- This isn't fun.
- This is expensive. And huge cost upfront.
- Focus on what your business does well



AWS By The Numbers



16 Regions



69 Availability Zones

165+ Services



Amazon Aurora



Amazon Managed Blockchain



Amazon EMR



Amazon GameLift



AWS Lambda



AWS IoT Button



AWS Config



Amazon Simple Notification Service



Amazon Sumerian



Amazon Chime



Amazon Connect



Amazon AppStream 2.0



Amazon Comprehend



AWS Application Discovery Service

Great Power, Great Responsibility

- There's hundreds of services available and are launched regularly
- Endless possibilities can have endless risks – need to manage and understand these risks.
- Breakdown of previous roles – infrastructure architects and security specialists are doing more 'cross functional roles' in the Cloud world



Enter Detective and Preventative Controls

Cloud Controls Overview

Preventive

Controls which prevent a deviation from occurring and serves as the first line of defence against organisational impact

Examples:

- Service Control Policies
- Identity-based or resource-based policies
- S3 block public access
- Inline validation of resources prior to deployment via a CI/CD pipeline (e.g., cfn-nag)

Detective

Controls which detect a deviation and trigger monitoring events/alerts to teams responsible for deviation and response management

Examples:

- Anomaly detection based on policies (e.g., unencrypted volume)
- Advanced event correlation through SIEM
- Bitcoin mining on an EC2 instance
- SSH brute force attack

Detective Controls Example

A developer changes the policy on an S3 bucket so they can access the files from home.

They remember they can **change** the **S3 Bucket Policy** to make it **public**. Perfect!



S3 Bucket



Bucket Policy:
Public read/write access
No logging
No Encryption

Detective Controls Solution

Receives event for policy change to S3 Bucket

Event matches rule for "No Public S3 Buckets unless whitelisted"

Security Incident created and team alerted so they can investigate

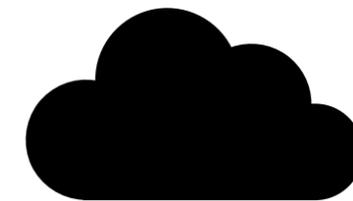
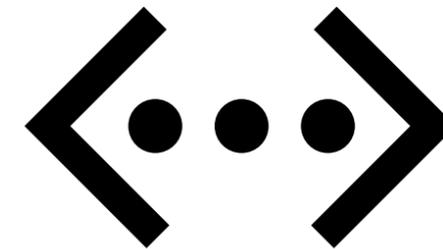
Problem Statement

Client requires a **Detective Controls Solution** to better **understand** and **manage** their **risk profile**

Selecting the Right Tool for the Job

Client Requirements

- A Detective controls framework that can scale with the organization
- Fit with existing development workflows
- Broad coverage of AWS services
- Needs to integrate with other tools and solutions for alerting, notification, reporting etc.
- An API for building out custom integrations, extending functionality



Aligning Controls to Standards

- Sourced team identified **218** Cloud Controls:
 - Aligned with industry controls frameworks:
 - CIS v1.2.0 - 05-23-2018
 - NIST 800-54 rev4
 - Provide high-level, provider neutral guidance on implementation
 - Cover broad level of categories (Access Management, Configuration Management ..)



Where We're Going, We Need
Automation.

Compliance as Code at Scale

We cannot manage rules, reports and alerts for 100+ AWS accounts

Define and maintain compliance configuration **as code**:

- ✓ Accounts
- ✓ Rules and profiles
- ✓ Notifications profiles
- ✓ Reports

Consistently apply **configuration baseline** across all environments and accounts.

Use **git workflows** for reviews

Protect 'master' release branch with:

- ✓ Mandatory reviews
- ✓ successful build requirement

Client Deliberates the Solutions....

Why Cloud Conformity?

Our client determined that Cloud Conformity best met their selection criteria and requirements

- Benefit of 450+ rules provided “out of the box”
- REST API with documentation on Github
- Works with many compliance standards
 - NIST
 - APRA
 - Singapore Monetary Authority
 - CIS
 - AWS Well Architected Framework



Trend Conformity?



All Cloud Accounts

Start typing...

- [Blurred list of cloud accounts]

Add an account
Create a new group

All accounts 149 AWS Accounts

Summary



98% Compliant
Good compliance level

Improve...

The latest Conformity Bot runs have performed 2184505 checks across all your accounts.

2136053 checks succeeded and 48452 failed.

Among failing checks, [] are **Very High** risk, [] are **High** risk, [] are **Medium** risk, [] are **Low** risk

Reports... Browse all checks...

Threat monitoring

Your accounts are not monitored.

User Activity

Open activity dashboard...

Event Monitoring

Open monitoring dashboard...

Account Identity

Event

Account Identity	Event
[Blurred data]	[Blurred data]
[Blurred data]	[Blurred data]
[Blurred data]	[Blurred data]

AMS01_ Settings

Summary

Jump to

- General s
- Rule setti
- Commun
- Access se
- Conform
- Real-time

General s

97

Good

Threat mo

The last activity on

User Activity

a minut

2 minute

Rule setti

Showing **472 rules**

Apply configurations from Profile...

Rule name	Categories	Service	Risk Level	Rule Id	
Security Group Port Range (Disabled)	Security	EC2	Low	EC2-001	Configure...
Unrestricted SSH Access	Security	EC2	Medium	EC2-002	Configure...
Unrestricted RDP Access	Security	EC2	Medium	EC2-003	Configure...
Unrestricted Oracle Access	Security	EC2	Medium	EC2-004	Configure...
Unrestricted MySQL Access	Security	EC2	Medium	EC2-005	Configure...
Unrestricted PostgreSQL Access	Security	EC2	Medium	EC2-006	Configure...
Unrestricted DNS Access	Security	EC2	Medium	EC2-007	Configure...
Unrestricted MsSQL Access	Security	EC2	Medium	EC2-008	Configure...
EC2-Classic Elastic IP Address Limit Updated	Performance-efficiency	EC2	Medium	EC2-009	Configure...
EC2-VPC Elastic IP Address Limit Updated (Disabled)	Performance-efficiency	EC2	Medium	EC2-010	Configure...
Account Instance Limit Updated	Performance-efficiency	EC2	Medium	EC2-011	Configure...
Security Group Excessive Counts					

But GUI's Don't Scale Well

- UI across 140+ accounts isn't feasible
- 140 accounts x 500 rules = 70 000 rules to configure (worst case)
- Each rule requires minimum 3 clicks to make a change
- 3 clicks x 500 rules x 140 accounts = 210 000 clicks
- Cloud Conformity Profiles feature addresses some of this

The Google logo is displayed in its standard multi-colored font (blue, red, yellow, green, red).

🔍 why do I have RSI

Google Search

I'm Feeling Lucky

Cloud Conformity Configuration as Code

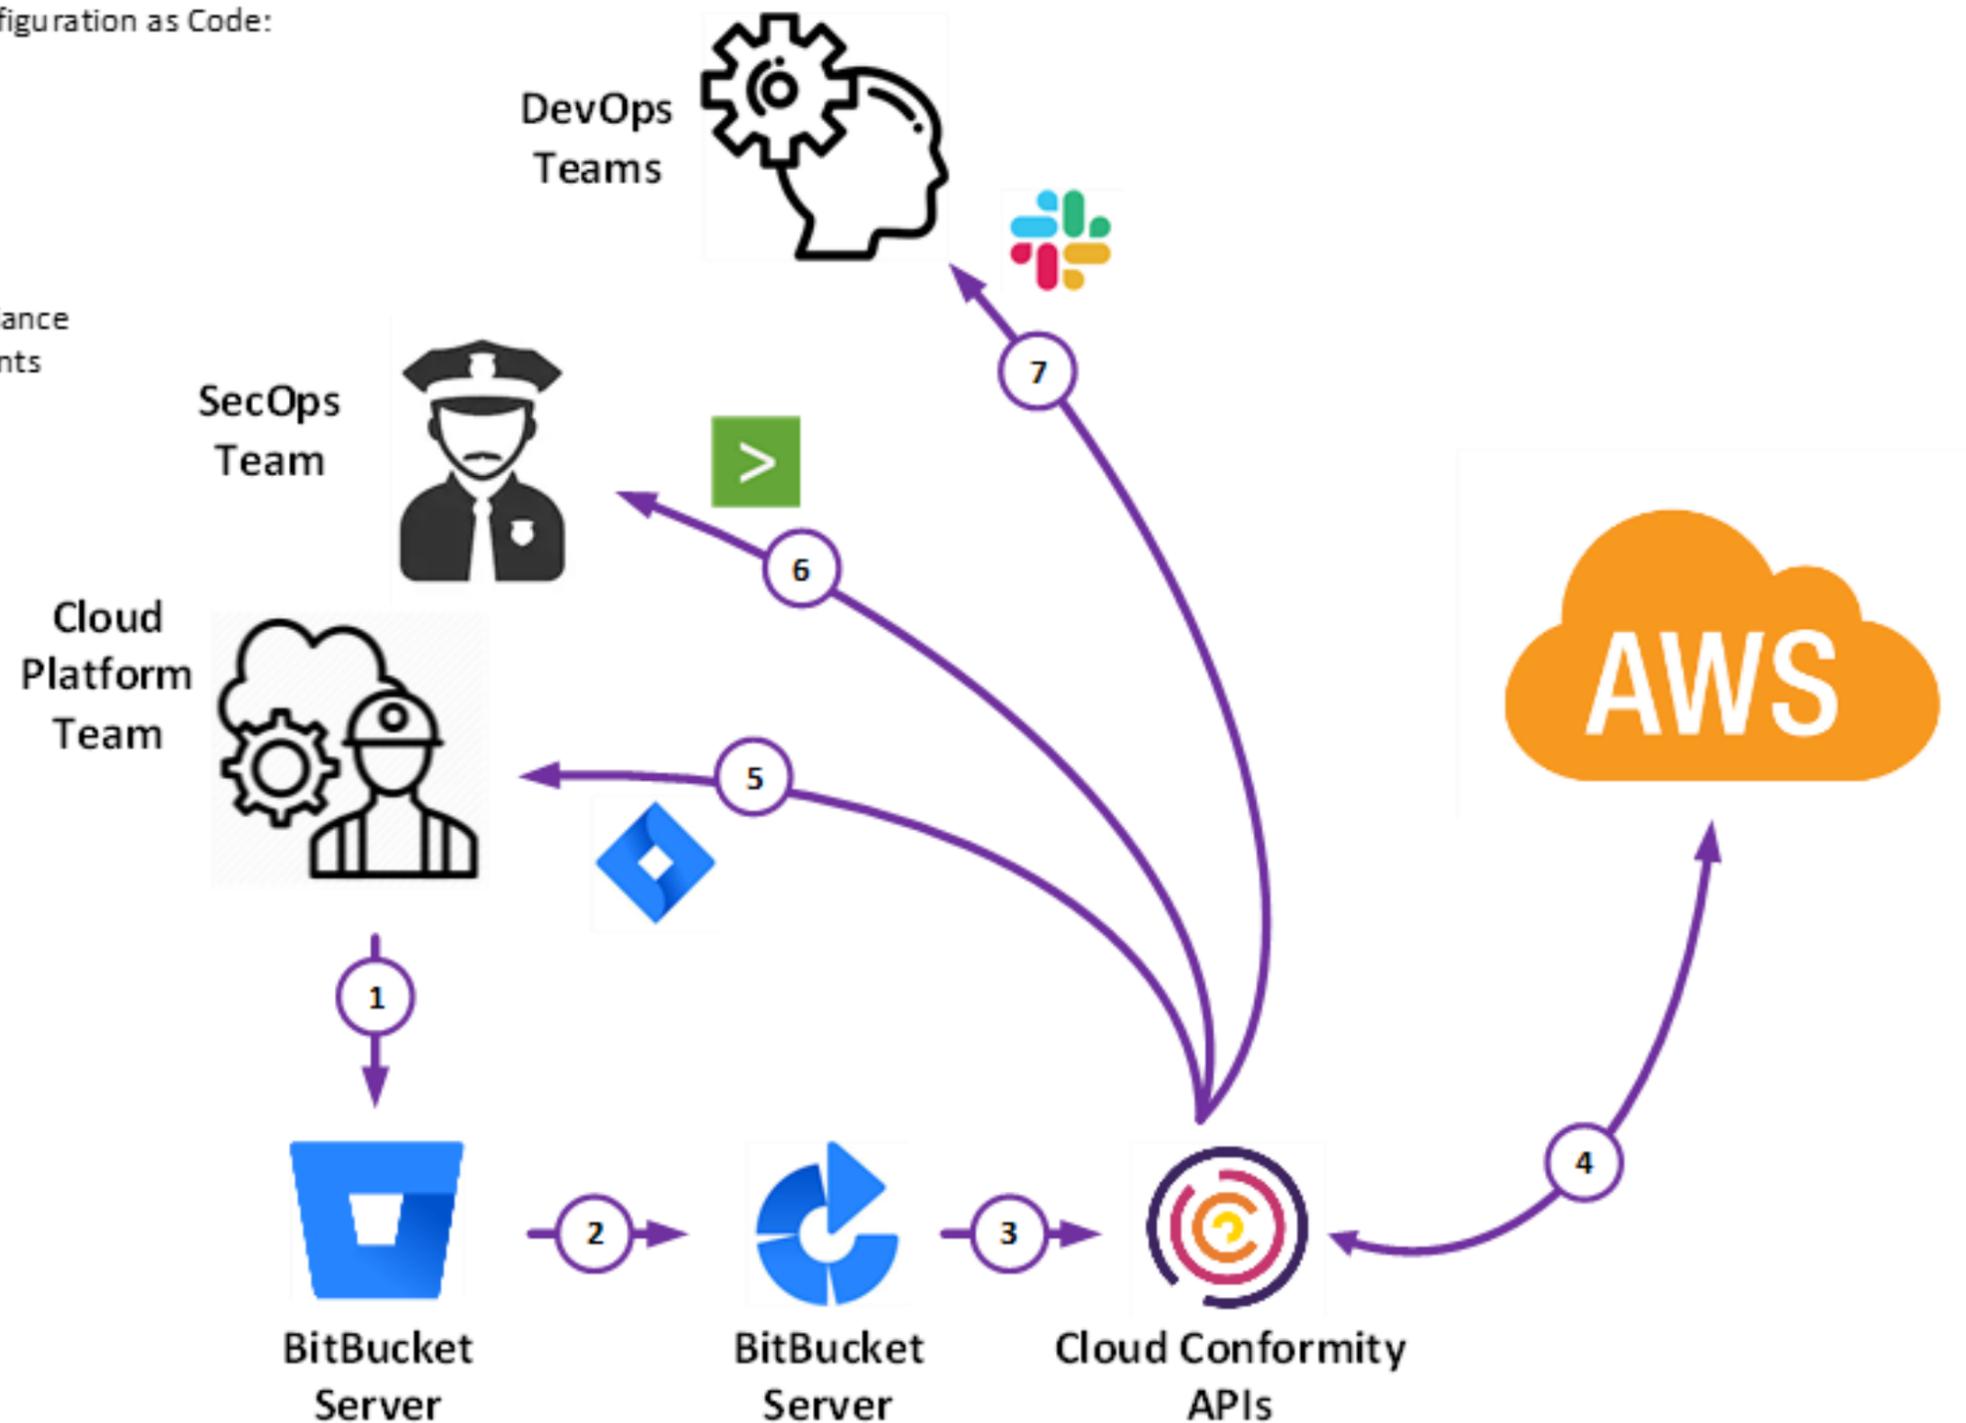
Key Features:

- Programmatically manage Cloud Conformity configuration at scale
- Applies **compliance-as-code** principles to Cloud Conformity workflows
- Allows for checking impact of changes with **-dry-run**
- Written in Python 3 with YAML for configuration files
- Jinja2 templates used to inject global variables across configuration
- Applies changes across all accounts in **parallel** for efficiency



Release Process

- 1 Commit Cloud Conformity configuration as Code:
 - Accounts
 - Rules
 - Notifications
 - Reports
- 2 Trigger deployment of Compliance configuration across all accounts
- 3 Apply configuration across All accounts in parallel
- 4 Trigger Ad-Hoc Scan Receive Real-Time results
- 5 Raise **JIRA** issues for HIGH RISK platform risks
- 6 Log all checks to **Splunk SEIM**.
- 7 Notify developers by **Slack** on application related checks



Using Python 3 to apply Detective Controls

- Python 3 is used to interact with the Cloud Conformity REST API
- Used Python 3 over Python 2 due to pending deprecation of Python 2 (you should be using Python 3)
- Concurrent Futures library used for multithreading the network requests so we can apply changes in parallel
- Can be run locally. Uses argparse module to parse CLI arguments
- Adapted to be used via AWS Lambda. No more servers.



Python – Benefits

- The team knew Python already. Faster to get started and be productive
- Dynamic typing gives us the ability to move quickly when prototyping
- *But also made some bugs harder to find*
- Handles JSON and REST API's *really well*
- Good library support
- Fast enough. Don't kid yourself that you need performance of C, C++ etc for everyday tasks

```
from conformity.logger import LOG
```

```
DEFAULT_BASE_URL = 'https://ap-southeast-2-api.cloudconformity.com/v1/{}'
```

```
class ConformityApi:
```

```
    def __init__(self, api_key=None, base_url=None):
        if not api_key:
            raise Exception("CC_KEY environment variable must be set or else I'm going on")

        self.api_key = api_key
        self.base_url = base_url or DEFAULT_BASE_URL
        self.headers = {
            'Content-Type': 'application/vnd.api+json',
            'Authorization': 'ApiKey {}'.format(self.api_key)
        }

        self.session = requests.Session()
        self.session.headers.update(self.headers)

    def get_user_email(self, user_id):
        ''' Get a user's email address from their CloudConformity ID '''

        user_data = self.get_user(user_id)

        if not user_data:
            return None

        email = user_data['data']['attributes']['email']

        return email
```

Python – Pain Points

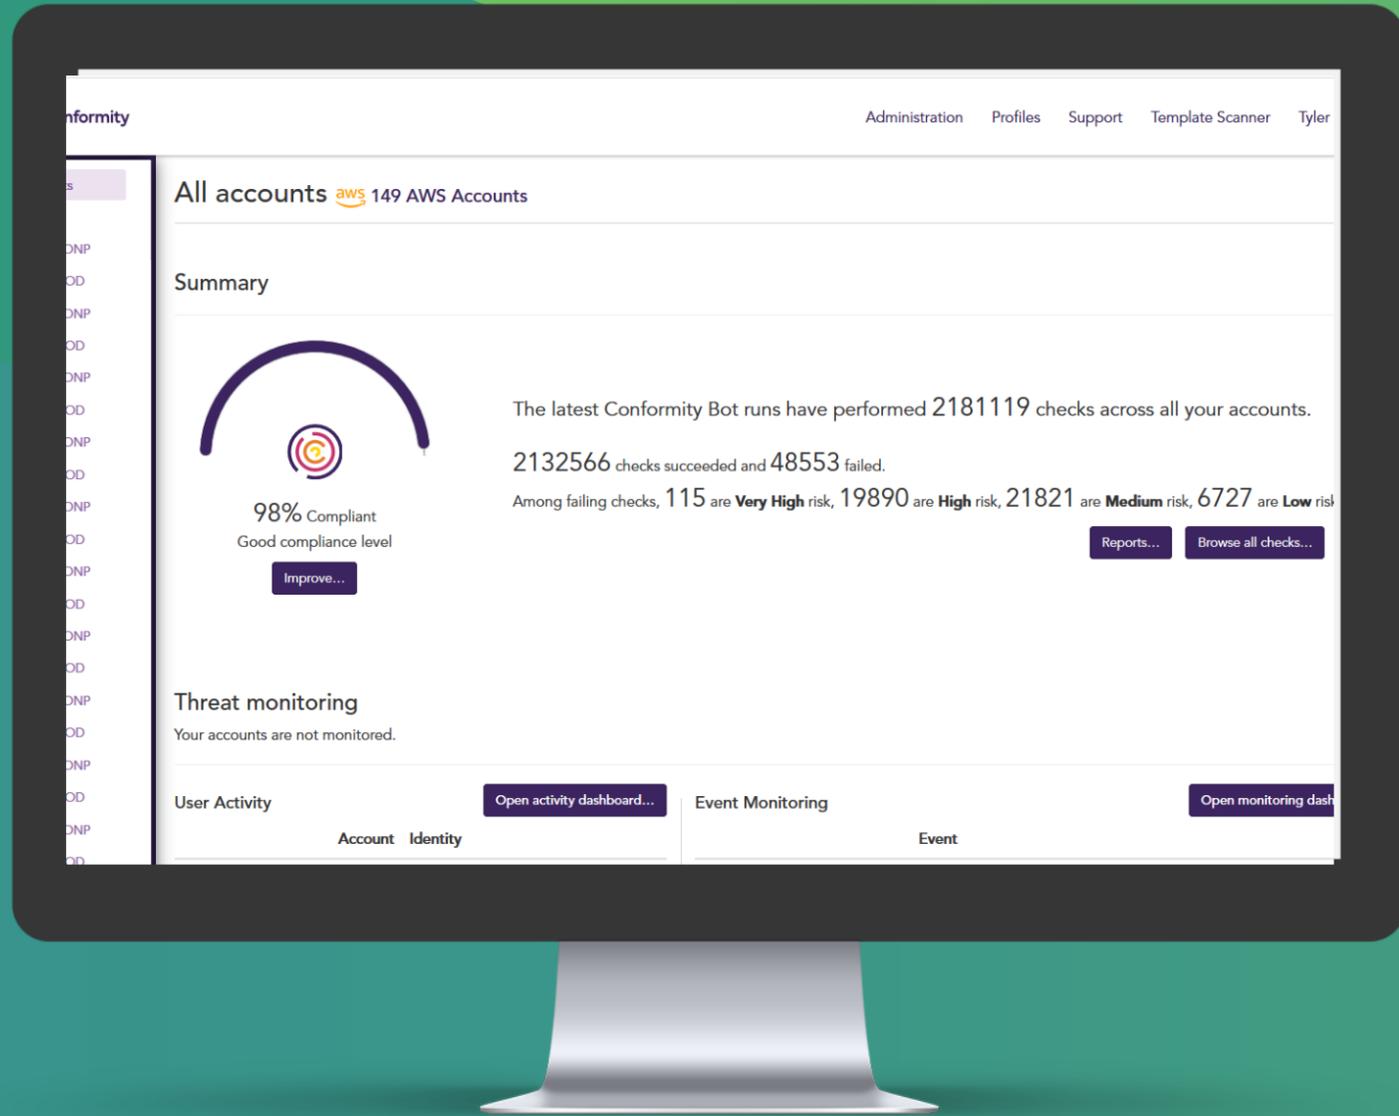
- Python 2/3 – some OS's still coming with Python 2
- Dependency management. Still not incredible.
- Dynamic typing can introduce fun bugs – not easy to detect
- Sometimes API's don't work how you might be expect
- Multithreading wasn't as easy as it should be. Concurrent Futures was the best solution.



```
for i in range(100):  
    print("ssssss")
```

A single-threaded Python

Cloud Conformity by the numbers



40M Checks each week

Over 40 million checks run each week run via Cloud Conformity across 140+ AWS Accounts.

All driven via existing development workflows

Automated deployment of rule changes across all accounts in minutes

Customer Outcomes



Customer Outcomes

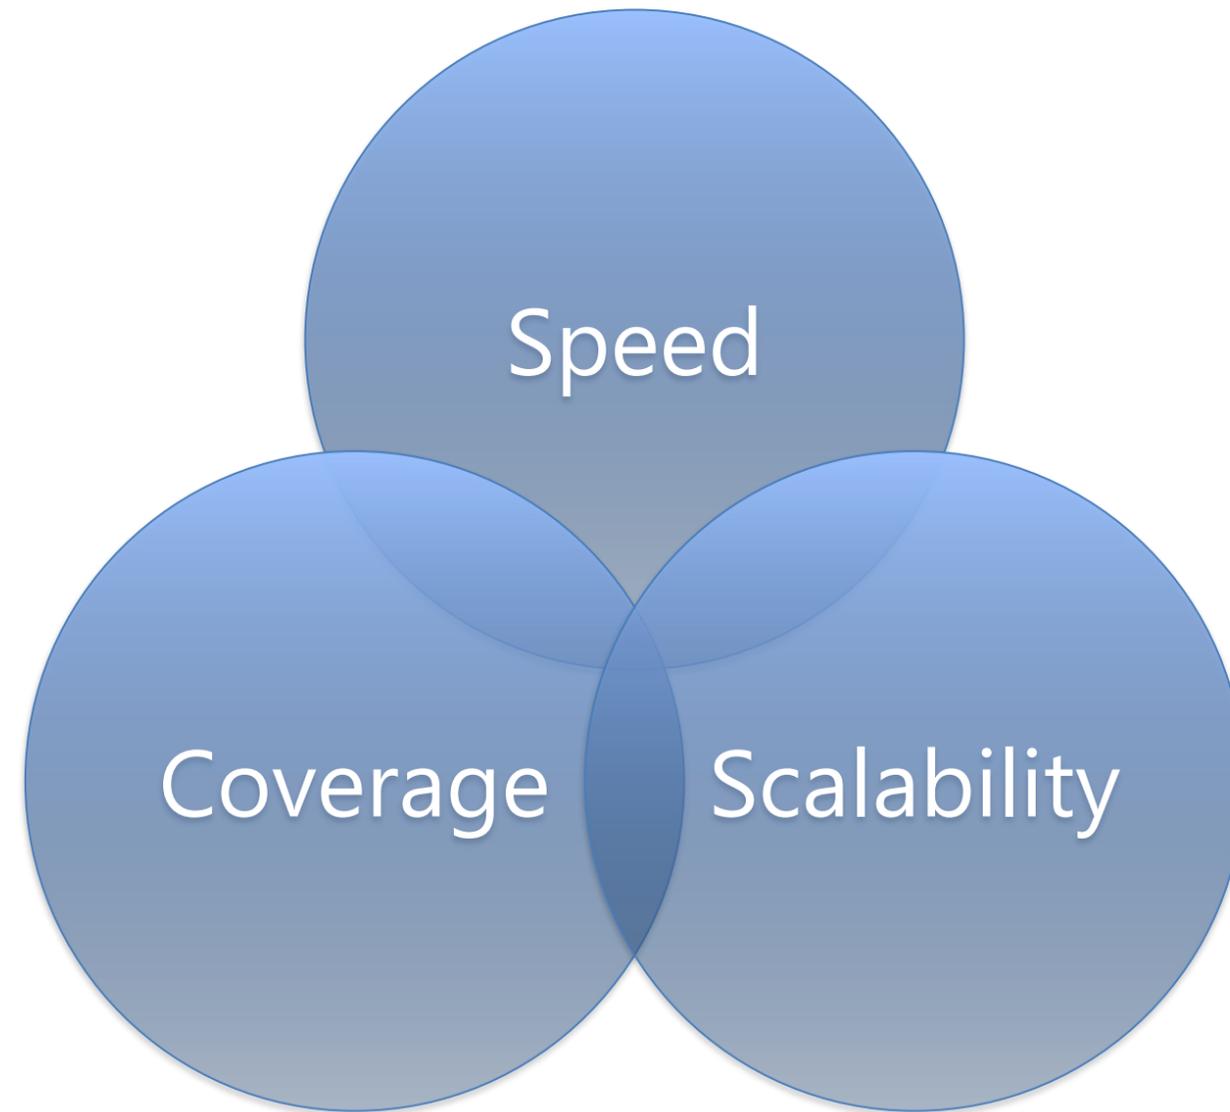


Scalability

Customer Outcomes

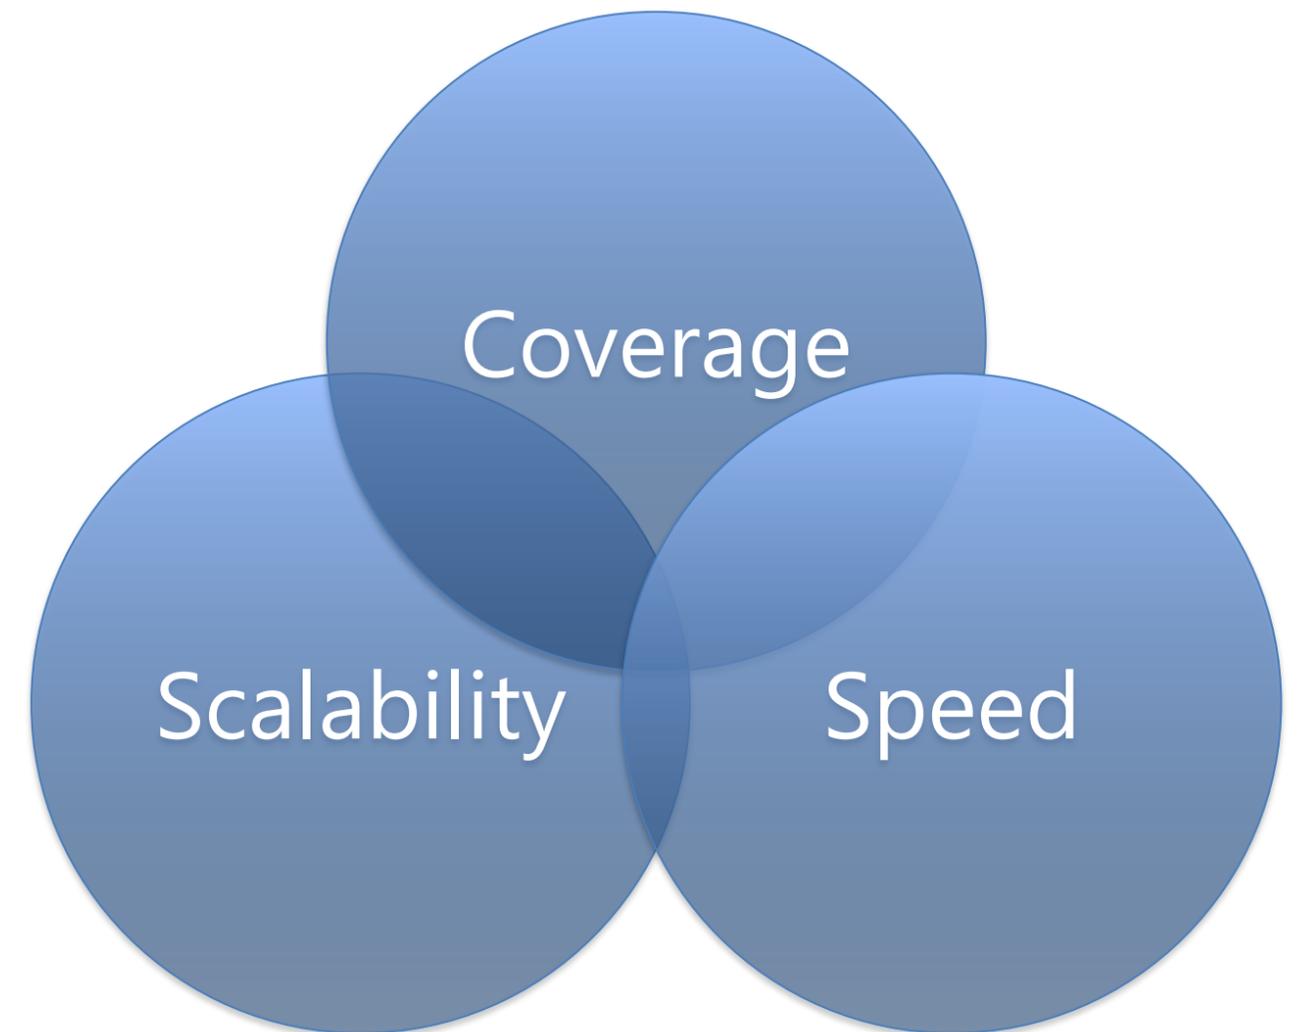


Customer Outcomes



Customer Outcomes

- ✓ Update rules/configuration programmatically across accounts (CLI, REST etc)
- ✓ Manage rules/configuration as code (DevSecOps approach)
- ✓ Integrated with existing development process and tooling
- ✓ Need to compare changes to current state before they are applied to platform (dry-run)
- ✓ Audit history of changes made via Conformity API
- ✓ Everything is Code



Key Challenges

Key Challenges

- Ensure relevance by reducing noise.
- Surfacing relevant information to the relevant teams
- Only surface the important stuff
- Give people enough information to help themselves – empowerment is key
- Remember not every app is the same



Cloud Conformity



Looking Forward

What comes next?

- **Auto-remediation:** Use Python, AWS Lambda and Boto3 to auto-remediate things picked up via Cloud Conformity
- **Custom Checks:** Easy way to add custom checks to Conformity – need to work out what this looks like in practice
- **Integration:** Build integrations for other tools/services with Cloud Conformity to get more 'single pane of glass' use out of it. Trend Micro integrations?

Key Points

- Python is a great general purpose programming language
- AWS offers lots of possibilities. And risks.
- Detective controls become really important as your AWS landscape scales
- Move to a “Everything as code” approach
- Not every Australian has been on Bondi Rescue.



A Marriage Made in Heaven



+



+



=





A Marriage Made in Heaven

sourced